

SYBASE®

Troubleshooting: Error Messages Advanced
Resolutions

Adaptive Server® Enterprise

15.0

DOCUMENT ID: DC00729-01-1500-01

LAST REVISED: January 2008

Copyright © 1987-2008 by Sybase, Inc. All rights reserved.

This publication pertains to Sybase software and to any subsequent release until otherwise indicated in new editions or technical notes. Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement.

To order additional documents, U.S. and Canadian customers should call Customer Fulfillment at (800) 685-8225, fax (617) 229-9845.

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase trademarks can be viewed at the Sybase trademarks page at <http://www.sybase.com/detail?id=1011207>. Sybase and the marks listed are trademarks of Sybase, Inc. ® indicates registration in the United States of America.

Java and all Java-based marks are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Unicode and the Unicode Logo are registered trademarks of Unicode, Inc.

All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian agencies.

Sybase, Inc., One Sybase Drive, Dublin, CA 94568.

Contents

About This Book	xix	
CHAPTER 1	Introduction	1
	Understanding error message severity levels	1
	Notifying the System Administrator	2
	Variables in error message text.....	2
	Finding object names in error message text	3
	Finding the Adaptive Server version	3
	Understanding the field order in the version string	4
	Creating error messages.....	4
	Reporting errors	4
	Sybase Technical Support checklist.....	5
	Sending error log fragments	6
	Reproducing problems	7
CHAPTER 2	Parser Errors (100s, 7380).....	9
	Error 102	9
	Error 107	11
	Error 195	12
	Error 7380	14
CHAPTER 3	Sequencer Errors (200s, 7783, 7788, 11000s, 14200).....	15
	Error 207	16
	Error 208	16
	Error 213	18
	Error 216	19
	Error 225	21
	Error 226	22
	Error 229	23
	Error 232	24
	Error 233	25
	Error 241	27
	Error 247	29

Error 257	30
Error 259	32
Error 265	33
Error 266	35
Error 268	38
Error 277	39
Error 7783	40
Error 7788	41
Error 11018	42
Error 11068	43
Error 11072	44
Error 14200	44

CHAPTER 4	Query Processor Errors (300s, 400s, 500s).....	47
	Error 311	47
	Error 313	49
	Error 314	51
	Error 403	53
	Error 414	54
	Error 428	55
	Error 511	56
	Error 512	57
	Error 515	58
	Error 530	60
	Error 539	61
	Error 540	61
	Error 546	63
	Error 547	64
	Error 551	65
	Error 584	66

CHAPTER 5	Access Method Errors (600s)	69
	Error 601	69
	Error 603	71
	Error 605	72
	Error 611	75
	Error 614	76
	Error 622	80
	Error 623	81
	Error 624	83
	Error 625	86
	Error 629	89
	Error 631	91

Error 644	95
Error 678	96
Error 691	97
Error 692	98
Error 693	100
Error 694	101
Error 695	103
Error 696	105
Error 697	106

CHAPTER 6 Memory Manager Errors (700s) 107

Error 701	107
Error 702	110
Error 703	111
Error 706	113
Error 707	113
Error 709	114
Error 714	115
Error 715	116
Error 717	116
Error 718	117
Error 719	117
Error 720	118
Error 721	119

CHAPTER 7 Buffer Manager Errors (800s) 121

Error 803	121
Error 804	122
Error 806	123
Error 813	125
Error 820	127
Error 821	127
Error 822	129
Error 823	131
Error 832	132
Error 834	133
Error 835	135
Error 840	135
Error 842	136
Error 847	137
Error 849	138
Error 852	139
Error 855	141

	Error 861	142
	Error 863	143
	Error 881	143
	Error 890	144
CHAPTER 8	Open Database Manager Errors (900s).....	145
	Error 903	145
	Error 905	147
	Error 906	149
	Error 908	149
	Error 911	150
	Error 913	152
	Error 916	153
	Error 921	155
	Error 924	156
	Error 925	157
	Error 926	160
	Error 930	161
	Error 935	163
	Error 940	164
	Error 941	166
	Error 945	167
	Error 950	168
	Error 965	169
CHAPTER 9	Abstract Plan Manager Errors (1000s).....	171
	Error 1024	171
CHAPTER 10	Page Manager Errors (1100s)	173
	Error 1105	173
	Error 1108	183
	Error 1120	184
	Error 1124	185
	Error 1127	186
	Error 1129	188
	Error 1131	190
	Error 1133	191
	Error 1141	193
	Error 1142	195
	Error 1143	199
	Error 1159	200
	Error 1161	200

CHAPTER 11	Lock Manager Errors (1200s).....	203
	Error 1203	203
	Error 1204	204
	Error 1205	206
	Error 1243	208
	Error 1244	208
	Error 1249	209
	Error 1251	210
	Error 1265	211
	Error 1279	212
CHAPTER 12	Sort Manager Errors (1500s).....	215
	Error 1501	215
	Error 1505	217
	Error 1506	218
	Error 1508	219
	Error 1509	220
	Error 1510	221
	Error 1514	223
	Error 1530	224
	Error 1531	224
CHAPTER 13	Initialization Errors (1600s)	227
	Error 1601	227
	Error 1602	229
	Error 1603	230
	Error 1605	231
	Error 1608	234
	Error 1613	235
	Error 1621	236
	Error 1622	238
	Error 1623	238
CHAPTER 14	Create Utilities Errors	
	(1700 - 1900s, 2100s, 2700s, 12800s).....	241
	Error 1702	242
	Error 1708	242
	Error 1732	243
	Error 1803	244
	Error 1808	245
	Error 1809	247
	Error 1810	248

	Error 1813	248
	Error 1820	249
	Error 1902	250
	Error 1903	251
	Error 1916	252
	Error 1928	253
	Error 2110	254
	Error 2714	255
	Error 2729	256
	Error 2753	257
	Error 2762	258
	Error 12818	260
	Error 12842	261
	Error 12881	262
CHAPTER 15	High Availability Errors (2200s).....	263
	Error 2243	263
CHAPTER 16	Character Set Conversion Errors (2400s)	265
	Error 2401	265
	Error 2402	266
	Error 2409	268
CHAPTER 17	dbcc Errors (2500s, 7900s, 12000s, 15000s)	271
	Error 2501	272
	Error 2502	273
	Error 2503	275
	Error 2507	277
	Error 2509	279
	Error 2510	281
	Error 2511	283
	Error 2513	284
	Error 2514	285
	Error 2517	287
	Error 2520	289
	Error 2521	290
	Error 2524	294
	Error 2525	296
	Error 2529	299
	Error 2540	301
	Error 2546	305
	Error 2547	310

Error 2550	312
Error 2558	314
Error 2559	321
Error 2572	324
Error 2573	325
Error 2574	325
Error 2575	327
Error 2578	328
Error 2582	330
Error 2583	332
Error 2591	334
Error 2596	335
Error 7901	338
Error 7902	340
Error 7928	343
Error 7930	343
Error 7939	346
Error 7940	348
Error 7949	350
Error 7989	352
Error 9961	353
Error 12964	354
Error 15056	355
Error 15080	356
Error 15082	356
Error 15084	357

CHAPTER 18 insert Errors (2600s) 359

Error 2601	359
Error 2610	360
Error 2615	361
Error 2619	365
Error 2620	366
Error 2626	368
Error 2628	369
Error 2630	371
Error 2631	372

CHAPTER 19 Procedure Manager Errors (2800s) 373

Error 2805	373
Error 2806	374
Error 2811	375
Error 2812	376

Error 2824	378
Error 2835	379
Error 2844	380

CHAPTER 20 dump and load Errors (3000s - 3200s, 4200s - 4300s)..... 381

Error 3019	382
Error 3020	383
Error 3021	384
Error 3105	385
Error 3120	385
Error 3143	386
Error 3144	387
Error 3161	387
Error 3162	388
Error 3201	388
Error 3203	389
Error 3211	391
Error 3212	391
Error 3216	392
Error 3225	393
Error 3230	393
Error 3233	394
Error 3240	395
Error 4204	396
Error 4205	400
Error 4207	401
Error 4216	402
Error 4221	403
Error 4222	404
Error 4305	406
Error 4322	406

CHAPTER 21 commit and abort Errors (3300s) 409

Error 3301	409
Error 3307	410
Error 3321	411
Error 3322	412

CHAPTER 22 Recovery Errors (3400s, 12500s) 413

Error 3401	413
Error 3403	414
Error 3404	416

	Error 3414	416
	Error 3429	417
	Error 3434	419
	Error 3445	420
	Error 3446	421
	Error 3447	422
	Error 3452	423
	Error 3454	424
	Error 3462	425
	Error 3470	426
	Error 12546	427
	Error 12547	428
	Error 12551	429
	Error 12583	429
	Error 12587	430
CHAPTER 23	Exception Handling Errors (3600s)	431
	Error 3604	431
	Error 3621	432
	Error 3624	432
	Error 3626	433
	Error 3641	435
	Error 3642	436
	Error 3643	436
	Error 3644	437
	Error 3646	437
	Error 3648	438
CHAPTER 24	drop Errors (3700s).....	441
	Error 3701	441
	Error 3702	442
	Error 3703	443
CHAPTER 25	Transaction Errors (3900s)	445
	Error 3904	445
	Error 3905	447
	Error 3906	448
	Error 3908	450
	Error 3917	451
	Error 3949	453
CHAPTER 26	dataserver Errors (4000s).....	455

	Error 4001	455
	Error 4002	456
	Error 4020	459
CHAPTER 27	View Resolution Manager Errors (4400s)	461
	Error 4430	461
CHAPTER 28	truncate table Errors (4700s)	463
	Error 4716	463
CHAPTER 29	Bulk Copy Utility Errors (4800s)	465
	Error 4801	465
	Error 4806	466
CHAPTER 30	alter table Errors (4900s)	469
	Error 4951	469
	Error 4953	470
	Error 4954	471
	Error 4956	471
	Error 4964	472
CHAPTER 31	alter database Errors (5000s)	475
	Error 5006	475
	Error 5013	476
	Error 5018	479
	Error 5034	480
CHAPTER 32	Disk Errors (5100s)	483
	Error 5115	483
	Error 5123	485
	Error 5142	487
CHAPTER 33	Transaction Coordinator Errors (5600s)	489
	Error 5602	489
CHAPTER 34	Open Client Errors (5700s)	491
	Error 5701	491
	Error 5704	492

CHAPTER 35	Configuration Errors (5800s, 10800s)	493
	Error 5806	494
	Error 5808	494
	Error 5824	495
	Error 5846	497
	Error 5847	497
	Error 5848	498
	Error 5849	499
	Error 5850	500
	Error 5851	501
	Error 5852	502
	Error 5853	502
	Error 5854	503
	Error 5857	504
	Error 5859	505
	Error 5861	506
	Error 5863	507
	Error 5865	508
	Error 5866	509
	Error 5867	510
	Error 5868	511
	Error 5893	511
	Error 10841	512
	Error 10865	513
CHAPTER 36	Process Kill Errors (6100s)	515
	Error 6103	515
	Error 6107	516
CHAPTER 37	Timestamp Errors (6900s).....	519
	Error 6901	519
	Error 6902	520
CHAPTER 38	Text Manager Errors (7100s).....	523
	Error 7101	523
	Error 7105	524
	Error 7114	526
	Error 7130	527
	Error 7134	529
CHAPTER 39	Distributed Database Network Errors (7200s).....	531
	Error 7201	531

Error 7205	532
Error 7207	534
Error 7211	535
Error 7212	536
Error 7214	537
Error 7215	538
Error 7218	540
Error 7220	540
Error 7221	543
Error 7223	545
Error 7227	546
Error 7234	547
Error 7235	548

CHAPTER 40	Threshold Manager Errors (7400s)	551
	Error 7401	551
	Error 7402	553
	Error 7403	554
	Error 7404	555
	Error 7406	556
	Error 7407	557
	Error 7408	558
	Error 7409	559
	Error 7410	560
	Error 7411	561
	Error 7412	562
	Error 7413	563
	Error 7414	563
	Error 7415	564
	Error 7417	565
	Error 7418	566
	Error 7438	567

CHAPTER 41	Auditing Errors (7600s)	569
	Error 7618	569

CHAPTER 42	Remote Procedure Call Errors (8000s)	571
	Error 8006	571
	Error 8009	572

CHAPTER 43	Descriptor Manager Errors (8200s)	575
	Error 8201	575

	Error 8203	576
	Error 8204	577
	Error 8207	578
	Error 8210	579
	Error 8211	580
	Error 8219	580
	Error 8239	582
	Error 8241	582
CHAPTER 44	sysindexes Manager Errors (8400s).....	585
	Error 8402	585
	Error 8419	586
CHAPTER 45	Process Status Structure Errors (8600s).....	589
	Error 8601	589
CHAPTER 46	Site Buffer Manager Errors (8700s).....	591
	Error 8704	591
CHAPTER 47	Disk Manager Errors (9000s)	593
	Error 9004	593
CHAPTER 48	Replication Agent Manager Errors (9100s - 9200s)	595
	Error 9122	595
	Error 9287	596
	Error 9289	597
	Error 9290	599
CHAPTER 49	Component Integration Services Errors (11200s).....	601
	Error 11201	601
	Error 11202	601
	Error 11203	602
CHAPTER 50	Configurable Shared Memory Dump Manager (12000s, 12300s) 605	
	Error 12036	605
	Error 12326	606
	Error 12328	607
	Error 12329	607
	Error 12330	608
	Error 12333	609

	Error 12334	609
	Error 12335	610
CHAPTER 51	Fault Isolation Utility Errors (12700s)	613
	Error 12716	613
	Error 12717	614
CHAPTER 52	Partition (14100s) and Partition Key Manager (14300s) Errors	617
	Error 14100	617
	Error 14108	618
	Error 14120	619
	Error 14124	619
	Error 14126	620
	Error 14129	620
	Error 14308	621
	Error 14313	622
CHAPTER 53	mount database Errors (14500s)	623
	Error 14503	623
	Error 14504	624
	Error 14511	624
	Error 14513	625
	Error 14560	626
	Error 14561	626
CHAPTER 54	Anchor Manager Errors (14600s)	629
	Error 14605	629
	Error 14618	630
	Error 14619	630
	Error 14620	631
	Error 14621	631
	Error 14622	632
	Error 14623	632
CHAPTER 55	XML Utility Manager Errors (14900s)	635
	Error 14901	635
CHAPTER 56	Message Service Errors (15100s).....	637
	Error 15101	637
	Error 15105	638

CHAPTER 57	Web Services Errors (15200s)	639
	Error 15219	639
CHAPTER 58	Encryption Errors (15400s)	641
	Error 15402	641
	Error 15403	642
	Error 15412	642
	Error 15413	643
	Error 15417	643
CHAPTER 59	Statistics Errors (15600s)	645
	Error 15600	645
	Error 15601	646
CHAPTER 60	System Procedure Errors (17000s, 18000s)	647
	Error 17461	648
	Error 17715	649
	Error 17716	649
	Error 17737	650
	Error 17870	650
	Error 17871	651
	Error 17872	651
	Error 17873	652
	Error 17874	652
	Error 17875	653
	Error 17903	653
	Error 17904	654
	Error 17905	654
	Error 17906	654
	Error 17910	655
	Error 18031	655
	Error 18032	656
	Error 18033	657
	Error 18123	657
	Error 18750	658
CHAPTER 61	Kernel Errors	661
	Asynchronous I/O not available	662
	Buffer mismatch	663
	Could not create shared memory	664
	Current process infected	665
	dopen error	667

dstart I/O request repeatedly delayed	669
Failed to open device	672
Feature is not licensed	673
File already in use	674
Memory too fragmented	675
Memory usage in procedure headers	676
Network information message.....	677
No more alarms available	678
Open objects parameter may be too low	679
os_attach_region errors	680
os_create_region errors	681
Read/write errors.....	683
Server is unnamed	685
Stack guardword corrupted	686
time slice error.....	689
t_rcvconnect error	691
udunmirror errors	691
No more engines can be offlined	693

About This Book

This guide contains troubleshooting procedures for problems that you may encounter when using Sybase® Adaptive Server® Enterprise. The problems addressed here are those which the Sybase Technical Support staff hear about most often.

Note Adaptive Server Enterprise is referred to as Adaptive Server in the remainder of this book.

Audience

This guide is intended for:

- Sybase system and database administrators
- Sybase Technical Support contacts
- Developers of applications using Sybase software

This guide assumes that you are thoroughly familiar with the Sybase products. If you are unfamiliar with any of the procedures described in this guide, call Sybase Technical Support for assistance.

How to use this book

This book provides detailed procedures for recovery from the most common Adaptive Server errors. Most resolutions are for numbered errors listed in the Adaptive Server *server.loc* file; however, some resolutions are for errors that do not have numbers.

Related documents

The Adaptive Server Enterprise documentation set consists of the following:

- The release bulletin for your platform – contains last-minute information that was too late to be included in the books.

A more recent version of the release bulletin may be available on the World Wide Web. To check for critical product or document information that was added after the release of the product CD, use the Sybase Technical Library.

- The *Quick Installation Guides* (UNIX and Windows) or the regular installation guide for your platform – describe installation, upgrade, and configuration procedures for Adaptive Server Enterprise and related Sybase products.

-
- The configuration guide for your platform (Windows or UNIX) – provides instructions for performing specific configuration tasks for Adaptive Server.
 - *ASE Replicator User's Guide* – describes how to use the Adaptive Server Replicator feature of Adaptive Server to implement basic replication from a primary server to one or more remote Adaptive Servers.
 - *Component Integration Services User's Guide* – explains how to use the Adaptive Server Component Integration Services feature to connect remote Sybase and non-Sybase databases.
 - *Enhanced Full-Text Search Specialty Data Store User's Guide* – describes how to use the Full-Text Search feature with Verity to search Adaptive Server Enterprise data.
 - *Glossary* – defines technical terms used in the Adaptive Server documentation.
 - *Java in Adaptive Server Enterprise* – describes how to install and use Java classes as datatypes, functions, and stored procedures in the Adaptive Server database.
 - *Job Scheduler User's Guide* – provides instructions on how to install and configure, and create and schedule jobs on a local or remote Adaptive Server using the command line or a graphical user interface (GUI).
 - *Migration Guide* – describes how to migrate data from one version of Adaptive Server to another.
 - *Messaging Service User's Guide* – describes how to use Real Time Messaging Services to integrate TIBCO Java Message Service and IBM WebSphere MQ messaging services with all Adaptive Server database applications.
 - *Monitor Client Library Programmer's Guide* – describes how to write Monitor Client Library applications that access Adaptive Server performance data.
 - *Monitor Historical Server User's Guide* – describes how to use Historical Server to obtain performance information for SQL Server[®] and Adaptive Server.
 - *Monitor Server User's Guide* – describes how to use Monitor Server to obtain performance statistics from SQL Server and Adaptive Server.
 - *Performance and Tuning Guide* – is a series of four books that explains how to tune Adaptive Server for maximum performance:

- *Basics* – the basics for understanding and investigating performance questions in Adaptive Server.
- *Locking* – describes how the various locking schemas can be used for improving performance in Adaptive Server.
- *Optimizer and Abstract Plans* – describes how the optimizer processes queries and how abstract plans can be used to change some of the optimizer plans.
- *Monitoring and Analyzing* – explains how statistics are obtained and used for monitoring and optimizing performance.
- *Query Processor Adaptive Server Enterprise* – describes the Query Processor in Adaptive Server Enterprise version 15.0 and how it is used to optimize query processing in Adaptive Server.
- *Quick Reference Guide* – provides a comprehensive listing of the names and syntax for commands, functions, system procedures, extended system procedures, datatypes, and utilities in a pocket-sized book (regular size when viewed in PDF format).
- *Reference Manual* – is a series of four books that contains the following detailed Transact-SQL information:
 - *Building Blocks* – Transact-SQL datatypes, functions, global variables, expressions, identifiers, wildcards, and reserved words.
 - *Commands* – Transact-SQL commands.
 - *Procedures* – Transact-SQL system procedures, catalog stored procedures, system extended stored procedures, and dbcc stored procedures.
 - *Tables* – Transact-SQL system tables and dbcc tables.
- *System Administration Guide: Volume 1* and *Volume 2* – provide in-depth information about administering servers and databases. These guides include instructions and guidelines for managing physical resources, security, user and system databases, and specifying character conversion, international language, and sort order settings.
- *System Tables Diagram* – illustrates system tables and their entity relationships in a poster format. Full-size available only in print version; a compact version is available in PDF format.

-
- *Transact-SQL User's Guide* – documents Transact-SQL, the Sybase enhanced version of the relational database language. This manual serves as a textbook for beginning users of the database management system. This manual also contains descriptions of the pubs2 and pubs3 sample databases.
 - *Troubleshooting and Disaster Recovery* – describes how to troubleshoot various system disasters, provides procedures for recovering from various system database disasters, and describes a variety of other useful recovery tasks.
 - *Troubleshooting: Error Message Advanced Resolutions* – (this book) contains troubleshooting procedures for problems that you may encounter when using Sybase Adaptive Server Enterprise. The problems addressed here are those which the Sybase Technical Support staff hear about most often.
 - *Unified Agent and Agent Management Console* – Describes the Unified Agent, which provides runtime services to manage, monitor and control distributed Sybase resources.
 - *Users Guide Sybase Software Asset Management* – introduces you to Sybase asset management concepts and provides instructions for establishing and managing license files when you install or administer Sybase products that use Sybase Software Asset Management (SySAM) 2.0.
 - *Using Encrypted Columns in Adaptive Server* – describes the encrypted column feature included with Adaptive Server Enterprise versions 15.0 and later.
 - *Using Adaptive Server Distributed Transaction Management Features* – explains how to configure, use, and troubleshoot Adaptive Server DTM features in distributed transaction processing environments.
 - *Using Sybase Failover in a High Availability System* – provides instructions for using Sybase Failover to configure an Adaptive Server as a companion server in a high availability system.
 - *Utility Guide* – documents the Adaptive Server utility programs, such as isql and bcp, which are executed at the operating system level.
 - *Web Services User's Guide* – explains how to configure, use, and troubleshoot Web Services for Adaptive Server.

- *What's New in Adaptive Server Enterprise?* – describes the new features in Adaptive Server Enterprise, the system changes added to support those features, and changes that may affect your existing applications.
- *XA Interface Integration Guide for CICS, Encina, and TUXEDO* – provides instructions for using the Sybase DTM XA interface with X/Open XA transaction managers.
- *XML Services* – describes the Sybase native XML processor and the Sybase Java-based XML support, introduces XML in the database, and documents the query and mapping functions that comprise XML Services.

Other sources of information

Use the Sybase Getting Started CD, the SyBooks™ CD, and the Sybase Product Manuals Web site to learn more about your product:

- The Getting Started CD contains release bulletins and installation guides in PDF format, and may also contain other documents or updated information not included on the SyBooks CD. It is included with your software. To read or print documents on the Getting Started CD, you need Adobe Acrobat Reader, which you can download at no charge from the Adobe Web site using a link provided on the CD.
- The SyBooks CD contains product manuals and is included with your software. The Eclipse-based SyBooks browser allows you to access the manuals in an easy-to-use, HTML-based format.

Some documentation may be provided in PDF format, which you can access through the PDF directory on the SyBooks CD. To read or print the PDF files, you need Adobe Acrobat Reader.

Refer to the *SyBooks Installation Guide* on the Getting Started CD, or the *README.txt* file on the SyBooks CD for instructions on installing and starting SyBooks.

- The Sybase Product Manuals Web site is an online version of the SyBooks CD that you can access using a standard Web browser. In addition to product manuals, you will find links to EBFs/Maintenance, Technical Documents, Case Management, Solved Cases, newsgroups, and the Sybase Developer Network.

To access the Sybase Product Manuals Web site, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

-
- The Sybase Developers Network (SDN) provides a variety of Web-based resources for developers who are interested in or already working with Sybase products. It also gives developers access to a community of peers who can provide assistance and advice. CodeExchange allows SDN members to freely exchange code samples, utilities, scripts and other tools.

To access the SDN Web site, go to SDN at
<http://www.sybase.com/developer/codexchange>

Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

❖ **Finding the latest information on product certifications**

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click Certification Report.
- 3 In the Certification Report filter select a product, platform, and timeframe, and click Go.
- 4 Click a Certification Report title to display the report.

❖ **Finding the latest information on component certifications**

- 1 Point your Web browser to Availability and Certification Reports at <http://certification.sybase.com/>.
- 2 Either select the product family and product under Search by Base Product; or select the platform and product under Search by Platform.
- 3 Select Search to display the availability and certification report for the selection.

❖ **Creating a personalized view of the Sybase Web site (including support pages)**

Set up a MySybase profile. MySybase is a free service that allows you to create a personalized view of Sybase Web pages.

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click MySybase and create a MySybase profile.

Sybase EBFs and software maintenance

❖ Finding the latest information on EBFs and software maintenance

- 1 Point your Web browser to the Sybase Support Page at <http://www.sybase.com/support>.
- 2 Select EBFs/Maintenance. If prompted, enter your MySybase user name and password.
- 3 Select a product.
- 4 Specify a time frame and click Go. A list of EBF/Maintenance releases is displayed.

Padlock icons indicate that you do not have download authorization for certain EBF/Maintenance releases because you are not registered as a Technical Support Contact. If you have not registered, but have valid information provided by your Sybase representative or through your support contract, click Edit Roles to add the “Technical Support Contact” role to your MySybase profile.

- 5 Click the Info icon to display the EBF/Maintenance report, or click the product description to download the software.

Conventions

The conventions for syntax statements in this manual are as follows:

Key	Definition
command	Command names, command option names, stored procedure names, utility names, utility flags, and other keywords are in bold.
<i><variable></i>	Variables, or words that stand for values that you fill in, are in italics and are normally surrounded with angle brackets <>. Do <i>not</i> include the brackets when typing in the value.
{ }	Curly braces indicate that you choose at least one of the enclosed options. Do not include braces in your option.
[]	Brackets mean choosing one or more of the enclosed options is optional. Do not include brackets in your option.
()	Parentheses are to be typed as part of the command.
	The vertical bar means you may select only one of the options shown.
,	The comma means you may choose as many of the options shown as you like, separating your choices with commas to be typed as part of the command.

Key	Definition
<>	Words in angle brackets should be replaced with the corresponding name (such as a table or device). Do not include brackets in your entry.

SQL syntax statements (displaying the syntax and options for a command) are printed as follows:

```
sp_dropdevice <device_name>
```

Examples showing the use of Transact-SQL commands are printed as follows:

```
1> select * from publishers
2> go
```

Examples of computer output are printed as follows:

```
pub_id  pub_name                city                state
-----  -
0736    New Age Books           Boston              MA
0877    Binnet & Hardley        Washington          DC
1389    Algodata Infosystems   Berkeley            CA
```

```
(3 rows affected)
```

Accessibility features

This document is available in an HTML version that is specialized for accessibility. You can navigate the HTML with an adaptive technology such as a screen reader, or view it with a screen enlarger.

Adaptive Server HTML documentation has been tested for compliance with U.S. government Section 508 Accessibility requirements. Documents that comply with Section 508 generally also meet non-U.S. accessibility guidelines, such as the World Wide Web Consortium (W3C) guidelines for Web sites.

Note You might need to configure your accessibility tool for optimal use. Some screen readers pronounce text based on its case; for example, they pronounce ALL UPPERCASE TEXT as initials, and MixedCase Text as words. You might find it helpful to configure your tool to announce syntax conventions. Consult the documentation for your tool.

For information about how Sybase supports accessibility, see Sybase Accessibility at <http://www.sybase.com/accessibility>. The Sybase Accessibility site includes links to information on Section 508 and W3C standards.

If you need help

Each Sybase installation that has purchased a support contract has one or more designated people who are authorized to contact Sybase Technical Support. If you cannot resolve a problem using the manuals or online help, please have the designated person contact Sybase Technical Support or the Sybase subsidiary in your area.



Introduction

This book contains detailed instructions on how to resolve the most frequently occurring Adaptive Server error messages. Most of the messages presented here contain error numbers (from the `master..sysmessages` table), but some error messages do not have error numbers, and occur only in the Adaptive Server error log.

Topic	Page
Understanding error message severity levels	1
Variables in error message text	2
Finding object names in error message text	3
Finding the Adaptive Server version	3
Creating error messages	4
Reporting errors	4

Understanding error message severity levels

Each numbered error message has an associated severity level that indicates the seriousness of the problem Adaptive Server has encountered.

Error messages with a severity lower than 19 display on the client. Errors with a severity of 19 and higher are sent to the error log. If the message indicates a fatal error, it typically includes a stack trace.

For more information about severity levels, see “Diagnosing System Problems” in the *System Administration Guide: Volume 1*.

Notifying the System Administrator

When errors with a severity level greater than 16 occur, inform the System Administrator. The System Administrator should monitor all errors that generate severity levels of 17 through 26, and is responsible for resolving these errors and tracking their frequency. To assist the System Administrator in resolving problems, print out a hard copy of the error log, which contains the stack trace from any fatal error.

Note Messages with a severity greater than 16, show severity 16 when the error is raised by `dbcc checktable` or `dbcc checkalloc`. This ensures that checks continue with the next object.

Variables in error message text

This table lists the symbolic constants that appear in error message text:

Symbol	Description
%d,%D	Decimal number
%x, %X, %.*x, %lx, %04x, %08lx	Hexadecimal number
%s	Null-terminated string
%.*s, %*s, %*.s	String, usually the name of a particular database object
%S_type	Sybase-defined structure
%c	Single character
%f	Floating-point number
%ld	Long decimal number
%lf	Double floating-point number

Finding object names in error message text

Many Adaptive Server error messages print only the logical page number, not the name of the affected table or index. The procedure “How to Find an Object Name from a Page Number” in the most recent version of *Troubleshooting and Disaster Recovery* explains how to find the table or index name using the logical page number supplied by the error message. You may want to review that section now so that you are familiar with the necessary procedures before an error occurs.

Warning! The procedure referenced in the preceding paragraph uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

Finding the Adaptive Server version

To find the version number of an Adaptive Server, use one of these binaries with a `-v` flag:

- `dataserver`, `diagserver`, `backupserver`, `isql`, `bcp`, and `buildmaster`

Note Only use `buildmaster` for version 12.0.x and earlier. `diagserver` is a diagnostic version of `dataserver`. Sybase Technical Support may ask you to run `diagserver` to help in problem solving.

- `defncopy`, and `langinstall`

UNIX example:

```
$$SYBASE/$SYBASE_ASE/bin/dataserver -v
```

```
Adaptive Server Enterprise/15.0/EBF 13443  
ESD#2/P/Sun_svr4/OS 5.8/ase150/2193/32-bit/FBO/Thu May  
18 03:21:41 2006
```

```
$$SYBASE/$SYBASE_ASE/bin/backupserver -v
```

Understanding the field order in the version string

The order in which fields appear in the `dataserver` version string depends on the version of the server you are running. Some examples follow.

- 1 Product.
- 2 Version number and SWR.
- 3 Release type: production (P), beta, or SWR version.
- 4 Platform identifier.
- 5 Operating system release under which the binary was compiled. This number is hard coded; it is not determined from your operating system.
- 6 Server.
- 7 Build number – a Sybase internal reference.
- 8 Mode – optimized (`dataserver`) or debug (`diagserver`).
- 9 Compilation date and time.

Creating error messages

You can create your own error numbers and messages based on Sybase error numbers. For more information, see “`sp_addmessage`” in the most recent version of the Adaptive Server guide *Reference Manual: Procedures*.

Reporting errors

When you report an error to Sybase Technical Support, have this information on hand:

- The environment in which the error occurred, including:
 - Adaptive Server version (see “Finding the Adaptive Server version” on page 3)
 - SWR (Software Release) number, if available
 - Client version

- Operating system version
- Hardware platform
- Production or development environment
- The message number, state, and complete error message text for any Adaptive Server errors. Include any numbers, database object types, or database object names that are included in the error message.
- The context in which the message was generated; specifically, what command was running at the time. Provide a copy of the error log, including all messages from just prior to Adaptive Server start-up to immediately before or immediately after the error message. Include the version number and date information printed at the beginning of the error log.
- The number of users who were logged into Adaptive Server when the error occurred.
- The frequency with which the error occurs and if the error can be reproduced at will.
- The client program that encountered the error (isql, Open Client™ CT-Library™, jConnect™ for JDBC™, and so on).

Sybase Technical Support checklist

This section lists the information you should collect to give to Sybase Technical Support before you call your Sybase Technical Support. Your primary site contact should have a copy of this information.

Providing Technical Support with this information will expedite your problem's resolution.

Table 1-1: Technical Support checklist

Item	Value
Site ID (customer number):	
Case Number:	
Priority:	
Support Engineer:	
Date Case Opened:	
Date Case Closed:	

Item	Value
Sybase Server version (include SWR [EBF] version if applicable):	
Operating system type and version:	
Front-end (client) product and version (include SWR [EBF] version if applicable):	
Client machine operating system type and version:	
Description of the problem (include error numbers, error messages, circumstances in which problem occurred):	
Other useful information (include copy of error logs, reproducible case, number of users on the system, average percent of CPU usage, any other pertinent information):	

Sending error log fragments

Whenever Technical Support asks you to send the error log to Sybase, include the log information from server start-up through the error message in question. Be sure to include the following:

- The lines just prior to start-up; for example:

```
kernel Adaptive Server Enterprise/12.5/SWR 9609 GA/P/Sun_svr4/OS 5.8
/main/1647/32-bit/FBO/Sat Jun  2 00:49:20 2001
```

- All the lines between start-up and the error; for example:

```
Error 2559, Severity Level 16: Data page number 0781 is empty.
```

- Any stack traces or unusual messages that occur near the error message in question.

For troubleshooting purposes, Sybase Technical Support may ask you to turn on trace flag 3601. When this trace flag is “on,” Adaptive Server writes a stack trace to the error log every time an Adaptive Server error is raised. Refer to “Starting Adaptive Server with trace flags” in the most recent version of the Adaptive Server Enterprise guide *Troubleshooting and Disaster Recovery*.

Reproducing problems

In general, resolution time in Sybase Technical Support may be reduced if you generate a reproducible test case demonstrating the error. Test cases can range from a single query that is not behaving as expected to a relatively complex code fragment that encounters a problem.

Whenever possible, test cases should be created using a small non-production database, such as the `pubs2` database. In many cases, this is a task that saves you from having to send large amounts of data to Sybase Technical Support.

Parser Errors (100s, 7380)

This section includes error messages for the Adaptive Server parser, which converts the text of the SQL statement to an internal representation called a query tree. See “Understanding Query Processing” in the *Performance and Tuning Series: Query Processing and Abstract Plans* for more information about parsers and query processing.

Error	Page
Error 102	9
Error 107	11
Error 195	12
Error 7380	14

Error 102

Severity 15

Message text Incorrect syntax near '%.*s'.

Explanation This error occurs when Adaptive Server detects a syntax error in a Transact-SQL[®] command or query. This error can occur when:

- A keyword is misspelled.
- A keyword or parameter is missing.
- The order of keywords is incorrect.
- You use an editor invoked from `isql` to write a SQL command or batch, end your SQL with “go” or another terminator, then run the file from `isql`. This causes an error only in some cases.
- A script that used to work no longer works after you changed the sort order or the character set of an Adaptive Server (for example, a script created on a server that originally was not case-sensitive and now Adaptive Server’s sort order has been changed to case-sensitive).

- A reserved word has been used in a query (for example, `select user = suser_name()` raises error 102 because “user” is a reserved word).
- You used a variable instead of a specific period of time following a `waitfor delay` statement. Variables are not supported in this case. For example, the following input would raise error 102:

```
1> declare @t char(9)
2> select @t = "00:00:01"
3> waitfor delay @t
4> go
```

- You used a variable for a database name and your query contains a parser error; error 102 is raised because the batch was never executed. Steps for executing a query are:
 - a Read the batch
 - b Parse the batch
 - c Optimize the batch
 - d Execute the batch

Any step that fails stops the query from processing.

Action

Check the spelling and syntax of the command specified in the error message. If it is wrong, correct it in your Transact-SQL statement and run it again.

You can invoke an editor such as `vi` from `isql` to edit a SQL statement or statements. However, for the statements to execute, you have to enter the command termination string (usually “go”) after you return to `isql`. If you include the command termination string in the editor, then save the file and return to `isql`, the statement will not execute. If you enter the command termination string again after returning to `isql`, you receive the following error (`vi` is the editor in this example; you may be using a different editor):

```
1> vi
1> select * from sysmessages where error = 102
2> go
3> go
Msg 102, Level 15, State 1:
Server 'REL1002_NAME', Line 2:
Incorrect syntax near 'go'.
```

If you are not sure whether a query contains a reserved word, use this query to see a complete list of reserved words:

```
1> select name from spt_values where type = "W"
2> go
```

Additional information	Refer to the most current version of the <i>Utility Guide, Transact-SQL User's Guide</i> , or the Adaptive Server Enterprise reference manuals for correct syntax.
Version in which this error is raised	All versions

Error 107

Severity	15
Message text	The column prefix '%.*s' does not match with a table name or alias name used in the query. Either the table is not specified in the FROM clause or it has a correlation name which must be used instead.
Explanation	Tables are specified in the <code>from</code> clause of a query. When Adaptive Server parses a query prior to execution, the name by which it knows the table is in the <code>from</code> clause—the table name if given alone, or a correlation name (alias) if one was specified. For example:

```
1> select c.cityname from cities c
2> go
```

In the above query, the table is known as `c` for purposes of the query. Names specified in other clauses refer back to this name to determine what table is intended. Error 107 can be raised:

- If a column prefix intended as a table name does not match any table name specified in the `from` clause.
- If a query does not use correlation names consistently. In Adaptive Server, queries that include correlation names must conform to ANSI requirements. Statements that specify correlation names but do not use them consistently raise error 107.

These restrictions apply to views as well as real database tables.

Action	Check your query for these errors and take corrective action:
--------	---

- 1 Check for a typing error; for example:

```
1> select * from titles
2> where title.code = 205
3> go
```

```
Msg 107, Level 15, State 1:
Server 'mfg', Line 2:
```

The column prefix 'title' does not match with a table name or alias name used

The column name in the where clause should be titles.code. This statement is also incorrect:

```
1> select t2.title_id from titles t1
```

The correct statement is:

```
1> select t1.title_id from titles t1
```

- 2 Ensure that correlation names are used correctly. For example, this statement is incorrect:

```
1> select title_id
2> from titles t
3> where titles.type = "french_cook"
```

The where clause can not use titles, because the from clause defines a correlation name for the table. The correct query is:

```
1> select title_id
2> from titles t
3> where t.type = "french_cook"
```

Additional information

The special case exists where a query that returns error 107 may report no error when the same type of correlation is used in a subquery. For example:

```
1> select * from mytable
2> where columnA =
3> (select min(columnB) from mytable m
4> where mytable.columnC = 10)
```

This query is a correlated subquery, and mytable.columnC refers to the outer table mytable. This query works because the same table is referred to in the inner and outer queries. In general, however, correlated subqueries can also generate error 107 when correlation names are used incorrectly.

Version in which this error is raised

All versions

Error 195

Severity

15

Message text

'%.*s' is not a recognized %S_MSG.

Explanation Error 195 is raised when a query contains illegal syntax. Specifically, it is raised when the Adaptive Server parser expects a specific syntax (such as a particular command, keyword, or punctuation mark), but instead encounters different syntax or missing text.

Error 195 is raised with these states:

State	Meaning
1	Invalid create procedure option.
2	Invalid timestamp keyword on writetext.
3	Invalid log keyword on writetext.
4	Invalid shutdown option, invalid trace keyword on kill, or invalid option on set lock.
5	Invalid set option; covers miscellaneous options that are set to on or off.
6	Invalid option on set offset.
7	All other invalid set options
8	Invalid option on set statistic.
15	Invalid quiesce database begin option
16	Invalid quiesce database end option

The following common query errors may raise a 195 error message:

- Using incorrect built-in function names (pre-12.5 version only; in Adaptive Server Enterprise 12.5 and later, a function that is not built-in is treated as an SQLJ function and raises error 14216 if it is not found).
- Using function names (pre-12.5 versions only) or using keywords that are not recognized by Transact-SQL.
- Using an invalid combination of punctuation marks, such as unpaired quotes.

Action Check the query syntax; the state of the error can be helpful in locating the source of the problem in a complex query.

Check the semantics to make sure that the keywords or options used in the query are valid where they are used in the query.

If you are running `xp_cmdshell`:

If the 195 error was raised when you attempted to run `xp_cmdshell`, check for invalid syntax and check the parameters supplied to the procedure. Combinations of single and double quotes can make it impossible for a command to be parsed; replace each double quote (") with two single quotes (') and retry.

If you are using Component Integration Services:

If the 195 error was raised when you queried an object in a remote server using Component Integration Services (CIS): your query may be using Transact-SQL syntax which is not understood by the (non-Sybase) remote server. To correct this, use the `sp_passthru` stored procedure, which allows you to pass a SQL statement using syntax native to the remote server. For example:

```
sp_passthru ORACLE, "select date from shiptable",
@errcode output, @errmsg output, @rowcount output,
@shipdate output
```

Additional information	See the <i>Reference Manual: Procedures</i> for more information about using <code>sp_passthru</code> .
Version in which this error is raised	All versions

Error 7380

Severity	10
Message text	WARNING: Line %d. Invalid value for '%S_MSG' specified. This value will be ignored, default value will be used instead.
Explanation	<p>You can specify a prefetch size on <code>select</code>, <code>update</code>, and <code>delete</code> statements for tables bound to caches with large I/O pools. Specify the prefetch size as the I/O size in kilobytes. Valid values for prefetch size are 2K, 4K, 8K, and 16K. The procedure <code>sp_helpcache</code> shows pools and their sizes for the cache to which an object is bound, or for the default cache.</p> <p>Error 7380 occurs when the prefetch size you specify is invalid. When this error occurs, Adaptive Server continues processing and the optimizer substitutes the “best” prefetch size available.</p>
Action	You must specify a prefetch size of 2K, 4K, 8K, or 16K. In addition, if the object is bound to a cache, make sure you use a valid prefetch size for that cache. Use <code>sp_helpcache</code> to see pools and their sizes for the cache to which the object is bound.
Additional information	See the Adaptive Server Enterprise reference manuals and performance and tuning guides for information about using prefetch sizes.
Version in which this error is raised	All versions

Sequencer Errors (200s, 7783, 7788, 11000s, 14200)

This section includes error messages for the Adaptive Server Sequencer. The Sequencer is the subsystem that controls command execution. The Sequencer handles the processing between steps in commands and also handles the procedural steps between commands such as if, goto, and while.

Error	Page
Error 207	16
Error 208	16
Error 213	18
Error 216	19
Error 225	21
Error 226	22
Error 229	23
Error 232	24
Error 233	25
Error 241	27
Error 247	29
Error 257	30
Error 259	32
Error 265	33
Error 266	35
Error 268	38
Error 277	39
Error 7783	40
Error 7788	41
Error 11018	42
Error 11068	43
Error 11072	44
Error 14200	44

Error 207

Severity	16
Message text	Invalid column name '%.*s'.
Explanation	<p>This error occurs when Adaptive Server tries to use a column name that does not exist. Adaptive Server supports delimited identifiers for table, view, and column names. Delimited identifiers are object names enclosed within double quotation marks. Using them allows you to avoid certain restrictions on object names. Delimited identifiers can begin with non-alphabetic characters, including characters that would not otherwise be allowed. Delimited identifiers may even be Transact-SQL reserved words.</p> <p>Delimited identifiers can cause some SQL statements to fail. For example, error 207 would be raised for the column “none” in this example:</p> <pre>1> select au_name, "none", au_fname 2> from pubs2..authors 3> go</pre>
Action	<p>Check the spelling of the column name specified in the error message. If the column name is wrong, correct it in your SQL statement and run it again.</p> <p>If this error message displays because of delimited identifiers, use single quotes instead of double quotes or turn delimited identifiers off. To turn delimited identifiers off, enter:</p> <pre>1> use <database_name> 2> go 1> set quoted_identifier off 2> go</pre> <p>When delimited identifiers is turned “on”, delimiting strings with double quotes causes Adaptive Server to treat them as identifiers.</p>
Additional information	Refer to the <i>Transact-SQL User's Guide</i> for information about delimited identifiers.
Version in which this error is raised	All versions

Error 208

Severity	16
----------	----

Message text %. *s not found. Specify owner.objectname or use sp_help to check whether the object exists (sp_help may produce lots of output).

Explanation This error occurs when Adaptive Server tries to use an object name that does not exist. If the object does exist, you may need to include the owner's name in the object name.

If the object is not owned by the user who is attempting to access it, and it is not owned by the database owner ("dbo"), all references to it must include the owner name. For example, if "user1" creates a table called test, all other users must prefix the table name with the owner name "user1" every time they refer to the object. The complete name in this example is user1.test. This requirement is based on Adaptive Server's naming convention for database objects:

```
<database>.<owner>.<object_name>.<column_name>
```

The default value for <database> is the current database and the default value for <owner> is the current user. Remember that the owner is part of the object name and it is therefore possible for two different users to have two different tables with the same <object_name> in the same database, for example, user1.test and user2.test. Refer to the *Transact-SQL User's Guide* for more details on naming conventions.

Temporary tables reside in tempdb and are automatically dropped when the user process or Adaptive Server connection that created them is exited. Furthermore, users cannot share temporary tables, even if they are created by the "dbo."

Including the fully qualified object name in an application may complicate the maintenance of the application. For example, if all references to a table include the database name, changing the database name could become quite difficult.

Action To resolve this error, refer to one of the following sections, depending on how much you know about the object in question.

If you do not know who owns the table

Use the sp_help procedure to display the owner (if the procedure is executed with no parameters), or query the system catalog to determine the object's owner and type. For example, to determine the name and type of the object table1, use the following query:

```
1> select owner = user_name(uid), name, type
2> from sysobjects where name = "table1"
3> go
```

If no rows are returned by this query, the object may reside in a different database or may not exist at all.

If you do not own the object in question

Avoid this error using either procedure:

- Include the owner's name in the object name. For example:

```
1> select * from user1.table1
2> go
```

- Have the “dbo” create the object. This allows any user to find the object without specifying an owner name.

If the object does not reside in the database

The error can be avoided by:

- Moving to the correct database with the `use <database_name>` command.
- Fully qualifying the object name with the database name. For example:

```
1> select * from database1.user1.table1
2> go
```

The owner name is not needed if you own the object or if it is owned by the “dbo.” For example:

```
1> select * from database1..table1
2> go
```

The appropriate permissions must also be set to allow access to this object. If these permissions are not provided, a 229 or 230 error results.

Version in which this error is raised

All versions

Error 213

Severity

16

Message text

Insert error: column name or number of supplied values does not match table definition.

Explanation	<p>This error occurs when there is a mismatch between the information specified in an insert statement and the definition of the table. Following are some examples of insert statements for the pubs2.titles table that cause this error (assuming that the titles table has nine columns):</p> <ul style="list-style-type: none"> • No column names are specified after the table name in an insert statement and there are two or more columns in the table. <pre> 1> insert titles values ('this is a test') 2> go Msg 213, Level 16, State 4: Line 1: Insert error: column name or number of supplied values does not match table definition.</pre> • No values are specified in an insert statement. <pre> 1> insert titles (title_id, title) values () 2> go Msg 213, Level 16, State 3: Line 1: Insert error: column name or number of supplied values does not match table definition.</pre>
Action	Check your insert statement to make sure you have specified column names and values for the columns.
Version in which this error is raised	All versions

Error 216

Severity	20
Message text	Attempt to automatically drop temporary table failed.
Explanation	<p>Adaptive Server goes through four phases while processing a query:</p> <ol style="list-style-type: none"> 1 Parsing – checks the query to make sure the syntax is correct. 2 Normalization – resolves the column and table names and generates a query tree. 3 Compilation – performs optimization and generates a query plan. 4 Execution – carries out the query plan.

During the first three phases of query processing, Adaptive Server creates two types of temporary tables:

- Temporary tables to aid in table resolution. This is necessary if the table will not exist until the command has been executed. For example:

```
1> create table x [information about table x]
2> select * from x
3> go
```

In this case, Adaptive Server creates a temporary copy of table *x* so it can resolve the `select` command.

- Other temporary tables prefaced with “#”, that exist for the life of the session (or the life of the procedure, if they are created in the procedure).

Error 216 occurs when Adaptive Server is unable to drop one of the temporary tables it created, probably caused by an Adaptive Server problem. Error 216 occurs with the following states:

State	Meaning
1	During the re-normalization phase, if Adaptive Server cannot drop a temporary table it created, error 216 occurs with state 1. These are definition-time temporary tables.
2	During the normalization phase, if Adaptive Server cannot drop a temporary table it created during the parsing phase, error 216 occurs with state 2.
3	During an abort of the normalization phase, if Adaptive Server cannot drop a temporary table, error 216 occurs with state 3.
4	When a process is killed, if Adaptive Server fails to clean up temporary tables, error 216 occurs with State 4. These can be definition-time or normal temporary tables.
5	If there is an error during the normalization phase, and Adaptive Server cannot drop a temporary table it created during normalization, error 216 occurs with state 5.
6	Similar to state 5.
7	During recompilation of a lightweight procedure, if Adaptive Server is unable to drop a definition-time temporary table, error 216 occurs with state 7.
9	If there is an error during creation of a temporary view, and Adaptive Server cannot drop a temporary table it created during normalization, error 216 occurs with state 9.
10	When processing nested SQL (for example, nested SQL statements within a Java method), if Adaptive Server cannot drop a temporary table at cleanup, error 216 occurs with state 10.
11	Similar to state 10.
12	When processing client or language cursors, if Adaptive Server cannot drop a temporary table at cleanup, error 216 occurs with state 12.
13	When regenerating a stored procedure during upgrade, if Adaptive Server cannot drop a normalization-time temporary table at cleanup, error 216 occurs with state 13.
14	Similar to state 13, but it applies to definition-time temporary tables.

Action	<p>Temporary tables left behind are not a problem in themselves. However, they do take up space in tempdb, and tempdb may eventually fill up. If you think this is a potential problem, shut down and restart Adaptive Server. This will re-create tempdb and drop all temporary tables.</p> <p>Error 216 can result from running out of locks. Check the Adaptive Server error log; if the 216 error is preceded by 1204 errors in the log, you may not have enough locks configured. See “Error 1204” on page 204 for instructions to correct that error.</p> <p>If error 216 persists, contact Sybase Technical Support.</p>
Version in which this error is raised	All versions

Error 225

Severity	11
Message text	Cannot run query--referenced object (name %.*s) dropped during query optimization.
Explanation	<p>This error occurs when Adaptive Server fails to build an execution plan for a stored procedure or an ad-hoc query because a referenced object was not found.</p> <p>Some reasons for this error include:</p> <ul style="list-style-type: none"> • The value of the number of open objects configuration parameter is too low • A problem exists in Adaptive Server <p>This is not a serious error and does not indicate corruption. However, it may prevent you from running some stored procedures and ad hoc queries.</p> <p>This error usually means a problem has occurred with Adaptive Server.</p>
Action	<p>Using sp_configure, review the value of the configuration parameter number of open objects and increase this value if feasible. For information about memory-related configuration parameters, refer to:</p> <ul style="list-style-type: none"> • The chapter “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> • Adaptive Server Enterprise performance and tuning guides <p>If you cannot increase the value of number of open objects or the 225 error occurs again, call Sybase Technical Support for assistance.</p>

Version in which this error is raised All versions

Error 226

Severity 16

Message text %s command not allowed within multi-statement transaction.

Explanation SQL commands are grouped into the following categories:

- SQL commands that are not allowed in transactions at all.
- SQL commands, such as data definition language (DDL) commands, that are allowed in transactions only if the required database option (`ddl in tran`) is set to `TRUE`.
- SQL commands that are allowed only if the transaction affects some other database. These commands include `create table`, `drop table`, and other commands that are run across databases to create or drop objects in another database when the database in which the objects are being created or dropped has the database option `ddl in tran` set to `TRUE`.

Error 226 occurs when Adaptive Server detects a command that is not allowed in a multi-statement transaction. A multi-statement transaction is a set of commands prefaced with the `begin transaction` command.

The following commands are never allowed in multi-statement transactions:

- `alter database`
- `create database`
- `dbcc reindex`, `dbcc fix_text`
- `disk init`
- `drop database`
- `dump database`, `dump transaction`
- `load database`, `load transaction`
- `select into`
- `set transaction isolation level`
- `truncate table`

- update statistics
- setuser

The following commands are not normally allowed in multi-statement transactions but you can use them if you use `sp_dboption` to set `ddl in tran` to `TRUE` first:

- create default, create index, create procedure, create rule, create schema, create table, create trigger, create view
- drop default, drop index, drop procedure, drop rule, drop table, drop trigger, drop view
- grant
- revoke

Action	<p>If the command is allowed in a multi-statement transaction when <code>ddl in tran</code> is set to <code>TRUE</code>, set <code>ddl in tran</code> to <code>TRUE</code> before running the transaction. Setting <code>ddl in tran</code> to <code>TRUE</code> causes locks on system tables, which can affect performance. You can check the current setting of <code>ddl in tran</code> with <code>sp_helpdb</code>.</p> <p>If the command is never allowed in a multi-statement transaction, execute it outside the multi-statement transaction.</p>
Additional information	<p>Some applications take SQL statements as input and run them for you. If the application uses <code>begin</code> and <code>commit</code> or <code>rollback</code> transaction to surround those statements, error 226 may occur. Refer to the user guide for your application to determine if this is the case.</p>
Version in which this error is raised	All versions

Error 229

Severity	14
Message text	<code>%s permission denied on object %S_OBJID, database %S_DBID, owner %.*s</code>
Explanation	<p>This error occurs when Adaptive Server attempts to access an object for which you do not have the appropriate permission.</p> <p>Error 229 can occur during any transaction which involves objects, such as <code>bcop</code>, <code>select</code>, or <code>update</code>.</p>

	<p>This error causes a rollback of the current transaction. Remaining commands in the batch are not executed and processing continues at the next batch.</p>
Action	<p>Check permission on the object named in the error message:</p> <pre>1> use <database_name> 2> go 1> sp_helprotect <object_name> 2> go</pre> <p>Where <i><database_name></i> is the name of the database in the error message and <i><object_name></i> is the name of the object in the error message.</p> <p>Grant the needed permission to the object (you either need to own the object or be the “sa” to do this). For example, to give Mary permission to insert into and delete from the titles table, type the following:</p> <pre>1> grant insert, delete 2> on titles 3> to mary 4> go</pre>
Additional information	<p>Refer to “Managing User Permissions” in the most recent version of the <i>Adaptive Server System Administration Guide: Volume 1</i> for detailed information about permissions.</p>
Version in which this error is raised	<p>All versions</p>

Error 232

Severity	16
Message text	Arithmetic overflow error for type %s, value = %f.
Explanation	<p>This error occurs when Adaptive Server detects an arithmetic overflow error during execution of a SQL statement. Arithmetic overflows occur when a variable has too few places to store the converted data. In particular, error 232 is raised when an overflow occurs for a variable of datatype float.</p> <p>Error 232 can occur during execution of the <code>convert</code> function. It can also occur during implicit conversion within Adaptive Server, either during a stored procedure run or during execution of a SQL statement which converts a float value to some other datatype.</p>

Action	<p>Determine which variable is causing the problem based on the datatype reported in the error message output and the SQL statement on which the error occurred. If the error occurs within a stored procedure, an additional error message displays showing the name of the stored procedure and the line number where the error occurred.</p> <p>Increase the size of the variable that is causing the error. If it is in a procedure, re-create the procedure.</p> <p>If this error occurs on a Sybase-supplied system stored procedure, contact Sybase Technical Support.</p>
Version in which this error is raised	All versions

Error 233

Severity	16
Message text	The column %.*s in table %.*s does not allow null values.%S_EED
Explanation	<p>During query processing, Adaptive Server performs an action called normalization, during which it resolves column and table names and generates a query tree. During normalization, when Adaptive Server prepares the target list for processing an insert or update, it checks to make sure nulls are allowed. Error 233 will be raised at compile time if, from within a stored procedure, you try to insert null values in a column which does not allow nulls.</p>

Error 233 is raised with the following states:

State	Meaning
1	Nulls are not allowed, the column is in the target list, and the value being inserted is null.
2	No default exists, nulls are not allowed, and the column is not in the target list.
3	For inserts to temporary tables, Adaptive Server checks to be sure all “not null” columns are listed explicitly. If a “not null” column is not listed in the insert statement, the error is raised with state 3.

Action	<p>When you create a table, you can explicitly define whether each column should allow null values. If you do not specify NULL or NOT NULL for a column when you create the table, the default value will be NOT NULL. If you use sp_dboption to set allow nulls by default to true for the database, any new table that is created will have columns with the default value NULL.</p>
--------	--

To determine whether a column has NULL or NOT NULL defined, enter commands like the following, replacing `pubs2` with the database in which your table resides and `titleauthor` with the table in which your column resides:

```
1> use pubs2
2> go

1> sp_help titleauthor
2> go
```

Name	Owner	Type							
titleauthor	dbo	user table							
Data_located_on_segment		When_created							
default		Oct 27 1994 10:09AM							

Column_name	Type	Length	Prec	Scale	Nulls	Default_name	Rule_name	Identity
au_id	id	11	NULL	NULL	0	NULL	NULL	0
title_id	tid	6	NULL	NULL	0	NULL	NULL	0
au_ord	tinyint	1	NULL	NULL	1	NULL	NULL	0
royaltyper	int	4	NULL	NULL	1	NULL	NULL	0

The `Nulls` column indicates whether null values are allowed. A value of zero (0) for the column means nulls are not allowed; a value of 1 means null values are allowed.

To change the default for a column from NOT NULL to NULL:

- 1 Use `bcp` to copy the data out of the existing table.
- 2 If you want the modified table to have the same name as the existing table, drop the old table.
- 3 Re-create the original table, specifying NULL for the column you want to change.
- 4 Use `bcp` to put back the data for the table.

If you want new tables in the database to have the default NULL, use the following commands for the database:

```
1> use master
2> go

1> sp_dboption <database_name>,
```

```

2> "allow nulls by default", true
3> go

1> use <database_name>
2> go

1> checkpoint
2> go

```

Where *<database_name>* is the name of the database whose behavior you want to change.

Version in which this error is raised

All versions

Error 241

Severity	16
Message text	Scale error during %S_MSG conversion of %s value '%s' to a %s field.
Explanation	<p>Error 241 occurs when Adaptive Server attempts to convert a given value from one numeric or decimal datatype to another, but is unable to do so because the scale of the receiving field (the “to” field in the message) is not large enough to accommodate the converted value.</p> <p>The error is raised during implicit conversions, which occur when Adaptive Server performs certain types of comparisons, and while the server stores and retrieves data. Adaptive Server automatically handles conversions between compatible datatypes, but may be unable to handle the numeric or decimal conversion due to the reasons mentioned above.</p> <p>For example:</p> <pre> 1> declare @numvar numeric(5,2) 2> select @numvar = 38.4434 3> go Msg 241, Level 16, State 2: Server 'vin_125_ocstest', Line 2: Scale error during implicit conversion of NUMERIC value '38.4434' to a NUMERIC field. </pre> <p>The scale of the numeric field <code>numvar</code> (2) is too small to accept the inserted value.</p>

Action

To resolve the error, perform any one of these options. Detailed instructions for each option follow this list.

- Set `arithabort numeric_truncation` to “off.”
- Use the `convert` function.
- Declare a large enough scale explicitly.

Setting `arithabort off`

Use `arithabort` to direct Adaptive Server not to abort the query due to a numeric truncation error. For example:

```
1> set arithabort numeric_truncation off
2> go
```

```
1> declare @numvar numeric(5,2)
2> select @numvar = 38.44*.72
3> select @numvar
4> go
```

```
(1 row affected)
-----
27.67
```

Using the `convert` function

Use the general-purpose `convert` function to specify the output. For example:

```
1> declare @numvar numeric(5,2)
2> select @numvar = 38.44
3> select @numvar=convert(numeric(5,2),@numvar*.72)
4> select @numvar
5> go
```

```
(1 row affected)
(1 row affected)
-----
27.67
(1 row affected)
```

Declaring a variable's scale

Declare a large enough scale when defining a variable. For example:

Instead of using:

```
1> declare @x decimal
2> select @x = 12.345
```


use:

```
1> declare @x decimal (5,3)
2> select @x = 12.345
```

Additional information The set arithabort and set arithignore options specify behavior following the loss of precision or (for numeric and decimal types) the loss of scale. For more information, see:

- The section “Arithmetic overflow and divide-by-zero errors” in the *Transact-SQL User’s Guide*
- The chapter “System and User-Defined Datatypes” in *Reference Manual: Building Blocks*.

When you write SQLJ Java methods (in Adaptive Server Enterprise version 12.5 and later), you can avoid 241 errors by ensuring that input parameters for the function calling the Java method use the proper scale and precision, as described in “Declaring a variable’s scale” on page 28.

Version in which this error is raised

All versions

Error 247

Severity

16

Message text

Arithmetic overflow during %S_MSG conversion of %s value '%s' to a %s field.

Explanation

Error 247 occurs when Adaptive Server attempts to convert a given value from one data type to another, but is unable to do so because of incompatibilities between the datatypes, or when the receiving field (the “to” field in the message) does not have a range large enough to accommodate the converted value.

The error may be raised during either explicit or implicit conversions. Explicit conversions occur when a query explicitly requests a conversion with the convert, intohex, or hextoint functions.

Implicit conversions occur when Adaptive Server is required to perform certain types of comparisons between heterogeneous datatypes, and when the server stores and retrieves data. Adaptive Server automatically handles many such conversions from one datatype to another, but may be unable to handle the conversion due to the reasons mentioned above.

For example:

```
1> create table deliver_dates
2> (itemno int,
3> dlydate smalldatetime)
4> go
```

```
1> insert into deliver_dates
2> values (42298, '12/12/2080')
3> go
```

Line 1: Arithmetic overflow during implicit conversion of VARCHAR value '12/12/2080' to a SMALLDATETIME field.

The date value being inserted is outside the range of SMALLDATETIME.

Action

Check the following problem areas when error 247 is raised:

- Is the range of the receiving datatype large enough to accommodate the converted value?
- Are the sending and receiving datatypes (the of and to datatypes in the message, respectively) compatible? For details, see the section “Datatype conversion functions” in the Transact-SQL Functions chapter in *Reference Manual: Building Blocks*.
- Is the query attempting to display a FLOAT value by converting it to numeric datatype? FLOAT values can only be displayed this way if no decimal digits are lost. Try increasing the precision of the numeric field so that this condition is met.
- If the query is updating a table, is there a trigger on the table that attempts to place the new value into another table? The receiving field in the second table may not have the correct range for the operation.

Additional information

Be sure to have the information listed in “Reporting errors” on page 4 when you contact Sybase Technical Support, including the text of the query that produced the error.

Version in which this error is raised

All versions

Error 257

Severity

16

Message text	Implicit conversion from datatype '%s' to '%s' is not allowed. Use the CONVERT function to run this query.
Explanation	<p>When performing certain types of comparisons, and while storing and retrieving data, Adaptive Server automatically handles many conversions from one datatype to another. These are called implicit conversions. You can explicitly request other conversions with the <code>convert</code>, <code>inttohex</code>, and <code>hexoint</code> functions.</p> <p>Error 257 is raised when Adaptive Server is unable to do an implicit conversion because of incompatibilities between datatypes.</p>
Action	<p>Corrective action depends on the exact situation in which the error occurred. Here are some typical situations where error 257 is raised, and how you can correct the problem:</p> <ul style="list-style-type: none">• If you attempt to perform comparisons on integer data with the <code>like</code> keyword; you must use the <code>convert</code> function on integer data if you want Adaptive Server to treat it as character data for comparisons.• If you are using embedded SQL, and Adaptive Server is unable to perform an implicit conversion between a database column and the corresponding host language variable. Consider using a different host variable to accomplish the conversion. For example, in a COBOL application, a column of money datatype requires a COMP-2 host variable.• If you attempt to insert quoted values for integer data. For example the following input raises error 257:<pre>1> create table citycodes 2> (cityname char(12), citycode smallint) 3> go 1> insert into citycodes 2> values ('Detroit', '123') 3> go</pre>Quotes are not allowed around the city code in this <code>insert</code> statement. <ul style="list-style-type: none">• When a statement implicitly exceeds the maximum length of a character datatype. For example, <code>columnX</code> is defined as <code>varchar(255)</code>. If you submit a query like this:<pre>1> select * from tableX 2> where columnX like '.....%'</pre>

and there are 255 characters preceding “%”, error 257 is raised. Adaptive Server treats the comparison string as if it contains more than 255 characters and assumes that it is a TEXT datatype, which it cannot implicitly convert.

Additional information For detailed information about datatype conversions, see the section “Datatype conversion functions” in the chapter “Using the Built-In Functions in Queries” in the most recent *Transact SQL User’s Guide*.

Before calling Technical Support, have the information on hand listed in “Reporting errors” on page 4, including the SQL statement that generated the error.

Version in which this error is raised All versions

Error 259

Severity 16

Message text Ad-hoc updates to system catalogs not enabled. A user with System Security Officer (SSO) role must reconfigure system to allow this.

Explanation System tables include all:

- Sybase-supplied tables in the master database; and
- Tables in the user databases that begin with “sys” and have an ID value in the sysobjects table of less than or equal to 100.

The system catalog contains the collection of system tables.

Error 259 occurs when you try to directly modify an Adaptive Server system table without first using sp_configure to set the allow updates configuration parameter. Setting this variable allows the System Administrator (“sa”) to make changes to system tables.

Warning! Serious problems can result from an incorrectly modified system table. Only change system tables when Sybase Technical Support or this manual instructs you to do so.

Action	<p>Do not use any of the procedures in this section before reading all warnings and cautions. Incorrect alterations of a system table can result in database corruption and data loss.</p> <p>To successfully modify the system catalog when directed by Technical Support or this guide, use this procedure:</p> <ol style="list-style-type: none"> 1 As “sa,” use <code>sp_configure</code> to allow changes to the system catalog: <pre style="margin-left: 40px;">1> sp_configure "allow updates", 1 2> go</pre> 2 Make changes to the system catalog as directed by Technical Support or elsewhere in this guide. Always use <code>begin transaction</code> when changing a system table, to protect against errors that could corrupt your database. 3 Immediately following the changes, use <code>sp_configure</code> to disallow subsequent changes to the system catalog, which is the default state for Adaptive Server: <pre style="margin-left: 40px;">1> sp_configure "allow updates", 0 2> go</pre> <hr/> <p style="text-align: center;">Warning! You must return Adaptive Server to its normal state, or anyone who can log in as “sa” can make changes to the system tables, or create stored procedures that alter system tables.</p> <hr/> <ol style="list-style-type: none"> 4 If you update <code>sysusages</code>, <code>sysobjects</code>, <code>sydatabases</code>, or <code>sysdevices</code>, restart Adaptive Server to update structures that are stored in cache and are related to those tables.
Additional information	See “allow updates to system tables” in the chapter “Configuring Parameters” in the <i>System Administration Guide: Volume 1</i> , and see the section on <code>sp_configure</code> in the <i>Reference Manual: Procedures</i> for more information about using (<code>sp_configure</code>) allow updates.
Version in which this error is raised	All versions

Error 265

Severity	16
Message text	Insufficient result space for %S_MSG conversion of %s value '%s' to a %s field.

Explanation	Error 265 can be raised when: <ul style="list-style-type: none">• Conversions to character data fails with error 265 if you attempt a conversion that would cause data truncation.• Certain system stored procedures are executed.
Action	This section describes scenarios under which error 265 occurs and procedures that you can use to resolve the error.

Errors in integer conversion

Conversions of integer to character data fail if the target format is not large enough to accommodate the data:

```
1> select convert (char(1), 500)
2> go
```

```
Msg 265, Level 16, State 1:
Server 'mfg1', Line 1:
Insufficient result space for explicit conversion of
INT value '500' to a CHAR field.
```

To correct this problem, choose a larger target format (char(3) in the example).

Errors in floating point conversion

Conversions of floating point to character data will fail if the target format is not large enough to accommodate the data:

```
1> select convert (char(10), 3.1415e)
2> go
```

```
Msg 265, Level 16, State 1:
Server 'mfg1', Line 1:
Insufficient result space for explicit conversion of
FLOAT value '3.1415000000000002' to a CHAR field.
```

In this example, the trailing part of the source value reported in the error message can be different from the entered value. This is because FLOAT is an approximate numeric datatype whose internal representation (and rounding upon display) are platform-dependent.

To correct this problem, choose a larger target format for display. Use the `str` function to determine the necessary format. The required format varies depending upon the number being converted and the accuracy of floating point numbers supported by your platform. To guarantee success, use a target of 25 characters.

Additional information	If loss of precision (rather than display format) is a concern in the application, consider using an exact numeric datatype such as integer, numeric or decimal.
Version in which this error is raised	Refer to the <i>Transact-SQL User's Guide</i> for details about the str function. All versions

Error 266

Severity	10
Message text	Transaction count after EXECUTE indicates that a COMMIT or ROLLBACK TRAN is missing. Previous count = %ld, Current count = %ld.
Explanation	<p>When a stored procedure runs, Adaptive Server maintains a count of open transactions, adding 1 to the count when a transaction begins, and subtracting 1 when a transaction commits. When you execute a stored procedure, Adaptive Server expects the transaction count to be the same before and after the stored procedure execution. Error 266 occurs when the transaction count is different after execution of a stored procedure than it was when the stored procedure began.</p> <p>Error 266 occurs most often when stored procedures are executed in chained mode. In chained mode, if there is no explicit begin transaction statement, an implicit begin transaction is issued. A begin transaction statement, whether implicit or explicit, increments the session transaction count by 1.</p> <p>In chained mode, the following statements cause an implicit begin transaction to be issued: select, update, delete, insert, and fetch. (The fetch command begins a transaction only when close on endtran is off in chained mode and cursors can remain open across transactions. See the <i>Reference Manual: Commands</i> for a discussion of the close on endtran option of the set command.)</p> <p>Each begin transaction statement in a stored procedure, whether implicit or explicit, must be balanced by a commit transaction statement, which decrements the session transaction count by 1. Otherwise, the transaction count remains higher when the stored procedure exits than it was at the beginning of execution.</p>

In addition, error 266 occurs when you are using nested procedures, and procedures at each level of nesting include `begin`, `commit`, and `rollback` transaction statements. If a procedure at a lower nest level opens a transaction and one of the called procedures issues a `rollback` transaction, error 266 occurs when you exit the nested procedure.

For example:

```
1> create procedure <proc1>
2> as
3> begin transaction
    .
    .
    execute <proc2>
4> commit transaction
5> go
1> create procedure <proc2>
2> as
3> begin transaction
    .
    .
4> rollback/commit transaction
5> go
```

If `<proc2>` executes a `rollback` transaction, error 266 is raised.

Action

In this example, execution of the stored procedure “test” results in error 266:

```
1> use pubs2
2> go

1> create procedure test as
2> select * from titles
3> go

1> sp_procxmode test, chained
2> go

1> set chained on
2> go

1> execute test
2> go
```

The transaction count is zero (0) before executing stored procedure “test”. When the stored procedure is executed, the `select` statement causes the transaction to begin implicitly. Since there is no balancing `commit` transaction for the implicit `begin transaction` the transaction count is not decremented in the stored procedure. The stored procedure returns to the caller with a transaction count of 1. A nonfatal 266 error is raised.

To balance the implicit begin transaction within a stored procedure, issue a commit statement before exiting the stored procedure. For example:

```
1> use pubs2
2> go

1> create procedure test as
2> select * from titles
3> commit transaction
4> go

1> sp_procxmode test, chained
2> go

1> set chained on
2> go

1> execute test
2> go
```

The transaction count is zero (0) before stored procedure execution. When the stored procedure is executed, the `select` results in an implicit begin transaction and the transaction count is incremented to one. The `commit transaction` in the stored procedure decrements the transaction count to zero (0). The transaction count is the same at the end of stored procedure execution as it was in the beginning.

You can use the `@@trancount` global variable to check the transaction count at the beginning and end of the procedure. If the count is not equal, commit or roll back transactions as appropriate.

Note It is safer to commit within the stored procedure than to open the transaction explicitly before executing the stored procedure. Opening the transaction before executing the stored procedure carries the risk of leaving a transaction open if the `execute procedure` statement fails and causing other problems as your procedures get more complicated.

Additional information

For more information on transactions within stored procedures and the `@@trancount` global variable, see “Checking the transaction nesting level with `@@trancount`” in the “Global Variables” section of the *Transact-SQL User's Guide*.

Version in which this error is raised

All versions

Error 268

Severity 16
Message text You can't run SELECT INTO in this database. Please check with the Database Owner.

Explanation This error occurs when you try to add rows to a table using the select into clause in a database that has the `sp_dboption` option `select into/bulkcopy` parameter disabled.

If this error occurs when you try to select into a temporary table, or while running a Sybase-supplied stored procedure, it is likely that the `select into/bulkcopy` option is not enabled in `tempdb`. Some stored procedures, for example, `sp_help` and `sp_helpsort`, trigger this error because they select into temporary tables (that all belong to `tempdb`) to get reports from the system tables, and to update them. If `tempdb` does not have the `sp_dboption` option `select into/bulkcopy` enabled, error 268 occurs.

Note When you install Adaptive Server, the `select into/bulkcopy` option is enabled in `tempdb` and turned “off” in all other databases.

You must enable the `sp_dboption` option `select into/bulkcopy` to perform nonlogged operations. For example:

- Perform a `select into` a permanent (rather than temporary) table.
- Perform a fast bulk copy with the `bcp` utility. Tables without triggers or indexes use the fast version of `bcp` by default, that is, they are not logged to save time.
- Execute the Transact-SQL command `writetext`. Using the `with log` option of `writetext` causes it to be logged, so the `select into/bulkcopy` parameter is unnecessary.
- Execute the DB-Library™ routine `dbwritetext`.

Warning! You cannot dump the transaction log to a device after performing a nonlogged operation. Always perform a dump database after the nonlogged operation is complete, since performing nonlogged operations leaves changes recorded in the transaction log unrecoverable.

Action Ask the System Administrator (“sa”) or database owner (“dbo”) to enable the `sp_dboption` option `select into/bulkcopy` for all databases affected by the query (including `tempdb` if applicable). The “sa” or “dbo” can use this procedure:

- 1 Use the master database and change the option:

```

1> use master
2> go

1> sp_dboption <database_name>,
2> "select into/bulkcopy", true
3> go

1> use <database_name>
2> go

1> checkpoint
2> go

```

- 2 To verify that the change is active, use sp_helpdb:

```

1> sp_helpdb <database_name>
2> go

```

In the output, the phrase `select into/bulkcopy` should appear, indicating that the option is enabled.

Warning! Dump your database before disabling the `select into/bulkcopy` option. If you insert nonlogged data (via a nonlogged operation) into the database, you will not be able to dump the transaction log to a device, and will not be able to recover your data.

By default, the `select into/bulkcopy` option is disabled when a database is first created. To change this default, use the procedure above to allow the option in the `model` database.

Additional information

If a table has indexes or triggers, then `bcp` will not run in the fast mode, so you do not need to set the `select into/bulkcopy` option to run `bcp`.

You cannot dump the transaction log to a device after performing a minimally logged operation. Trying to dump the transaction log after a minimally logged operation causes error messages that instruct you to dump the database instead.

Version in which this error is raised

All versions

Error 277

Severity

16

Message text	There was a transaction active when exiting the stored procedure '%.*s'. The temporary table '%.*s' was dropped in this transaction either explicitly or implicitly. This transaction has been aborted to prevent database corruption.
Explanation	<p>When exiting a stored procedure, Adaptive Server checks to see whether there is an active transaction (a transaction that has not been committed), then checks whether any temporary objects exist for the uncommitted transaction.</p> <p>If temporary objects exist, the transaction is rolled back, the temporary objects are dropped, and error 277 is raised.</p>
Action	<p>In stored procedures, make sure all begin transaction statements have corresponding commit transaction or rollback transaction statements.</p> <p>In chained mode, if there is no explicit begin transaction statement, an implicit begin transaction is issued. If you use chained mode, an explicit commit transaction or rollback transaction statement is required to end the transaction.</p> <p>If you confirm that all begin transaction statements have corresponding commit transaction or rollback transaction statements, check to see whether the stored procedure is exiting without completing its processing.</p>
Version in which this error is raised	All versions

Error 7783

Severity	10
Message text	Internal error -- server failed to do garbage collection on this procedure, id = %ld. Please save the query tree and procedure text and inform the technical support. This is not a critical error, so server continues processing.
Explanation	<p>Remapping a stored procedure, view or trigger generally causes the object to grow. After remapping, Adaptive Server compresses the object. Error 7783 occurs when Adaptive Server is unable to reclaim space after remapping the query tree. This occurs during execution of:</p> <ul style="list-style-type: none">• Stored procedures or triggers• Recovery, usually following a load database

This error is an Adaptive Server problem and is frequently accompanied by error 7961 (“Remapping utility - Wrong kind of node passed to tree remapping. This is an internal system error.”).

Although the server is unable to reclaim space, error 7783 does not cause severe problems and the procedure or trigger continues processing.

Action

Take these steps to resolve the error:

- 1 Identify the stored procedure, view, or trigger:

```
1> select object_name(<object_id>)
2> go
```

Where *<object_id>* is the ID named in the error message.

- 2 Check that the object text is available. It is important to verify this before you continue with the next step.
- 3 Drop and re-create the object. Note that you are re-creating the stored procedure, view, or trigger, and no changes are needed to the underlying tables.
- 4 If the error results in a stack trace or is accompanied by the 7961 error, you may need to upgrade to a version in which the problem is resolved. Call Sybase Technical Support for assistance.

Additional information

Have the information listed in “Reporting errors” on page 4 before calling Sybase Technical Support, including the output from `sp_configure`, `sp_who`, and `dbcc lock`.

Version in which this error is raised

All versions

Error 7788

Severity

10

Message text

```
WARNING: An explicit or implicit HOLDLOCK or NOHOLDLOCK
was specified for table '%.*s'; however, since this is
an isolation level 0 scan, the HOLDLOCK or NOHOLDLOCK
is ignored.
```

Explanation	<p>Isolation level zero (0) allows transactions to read uncommitted data (dirty reads). You can specify isolation level 0 for your queries using the <code>at isolation</code> syntax, or you can specify level 0 for your session as part of the <code>transaction isolation level</code> option of the <code>set</code> command. Adaptive Server's default transaction isolation level is 1.</p> <p>The <code>holdlock</code> parameter makes a shared lock on a specified table or view more restrictive by holding it until the completion of a transaction (instead of releasing the lock as soon as the required data is no longer needed, whether or not the transaction has been completed). Setting <code>transaction isolation level 3</code> with a <code>set</code> command implicitly applies a <code>holdlock</code>.</p> <p>The <code>noholdlock</code> parameter prevents Adaptive Server from holding any locks acquired, regardless of the transaction level currently in effect.</p> <p>Error 7788 is a warning message that is raised during query execution when level 0 (zero) scans are active and the query includes an explicit or implicit <code>holdlock</code> or <code>noholdlock</code> parameter. The query is executed and the <code>transaction isolation level 0</code> option of the <code>set</code> command takes precedence over the <code>holdlock</code> or <code>noholdlock</code> parameter of the query.</p>
Action	<p>To use <code>holdlock</code> or <code>noholdlock</code> in your query, use the <code>transaction isolation level</code> option of the <code>set</code> command to change the isolation level to a value other than 0 (zero) for your session:</p> <pre>1> set transaction isolation level <n> 2> go</pre> <p>Where <code><n></code> is the value to which you want to set the isolation level. Then re-enter your query.</p>
Additional information	<p>Refer to the <i>Transact-SQL User's Guide</i> for information about selecting an isolation mode for a query or for your session.</p> <p>Refer to the <i>Reference Manual: Commands</i> for information about <code>holdlock</code> and <code>noholdlock</code>.</p>
Version in which this error is raised	All versions

Error 11018

Severity	10
Message text	XP Server must be up for ESP to execute.

Explanation	XP Server is an Open Server application that runs on the same machine as Adaptive Server and lets you implement extended stored procedures (ESPs). Error 11018 is raised when Adaptive Server is unable to contact XP Server to execute an ESP.
Action	<p>Verify that your <i>interfaces</i> file (UNIX) or <i>sql.ini</i> file (Windows) has the correct entry for XP Server. The XP Server name must be all uppercase and should have the format <i>ASE_SERVERNAME_XP</i>.</p> <p>Also verify that the server name appears in <i>syssservers</i>, using the <i>isql</i> command:</p> <pre>1> select srvname from master..syssservers 2> go</pre> <p>If you do not find an XP Server entry or the entry is incorrect, refer to your platform's installation and configuration guide for details about installing XP Server.</p>
Additional information	Related information appears in the section "Solving common XP Server problems" in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	All versions

Error 11068

Severity	20
Message text	Transaction was found in the incorrect state of '%1!'. The expected state was '%2!'
Explanation	Adaptive Server is performing a generic cleanup after executing a statement. While checking to see if updates associated with the transaction are complete, Adaptive Server found a transaction in the wrong state.
Action	This error is raised most frequently when handling parallel processing, when resources are low. To avoid this error, increase the memory per worker process. Occasionally this error is preceded by another error (or shows a previous error in the error log) that reports you have run out of some configurable resource. If that is the case, address that issue first before trying to resolve error 11068.

Additional information	If the error persists, a configurable shared memory dump (CSMD) may be helpful in determining the root cause. Information about CSMD can be found in the chapter “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.0.0.7 and later

Error 11072

Severity	18
Message text	An internal conversion error was found while converting datatype `%s` to datatype `%s`.
Explanation	Adaptive Server received an error while running the internal routine that converts datatypes.
Action	Check the displayed datatypes to verify if there should be a legal conversion between them. See “Datatype Conversion Functions” in <i>Reference Manual: Building Blocks</i> .

Note Some Java datatypes do not have an SQL equivalent. For more information, see “Datatype mapping between Java and SQL” in the “Reference topics” chapter of the *Java in Adaptive Server Enterprise* manual for valid mappings. In some circumstances, attempting to convert a Java datatype can result in an 11072 error being raised.

Additional information	If the error persists, contact Sybase Technical Support, and have on hand the SQL text that generated this error.
Version in which this error is raised	Version 12.5.0.3 and later

Error 14200

Severity	16
Message text	The specified java signature is invalid.

Explanation	<p>SQLJ functions provide a layer on top of Java methods. They are used to invoke Java methods as SQL functions in Adaptive Server.</p> <p>The SQLJ function specifies the datatypes of the function's parameters, which are collectively known as the SQLJ function signature. You can also specify the datatypes of the underlying Java method's parameters, collectively known as the Java method signature.</p> <p>Error 14200 is raised when the SQLJ function explicitly specifies the Java method signature, but this signature does not map to the SQLJ function signature.</p> <p>Consider the following Java routine that uses a Java integer:</p> <pre>public class Jobcode { public static String job2(Integer jc) { if (jc == null) return null; else if (jc.intValue() == 1) return "Admin"; else if (jc.intValue() == 2) return "Sales"; else if (jc.intValue() == 3) return "Clerk"; else return "unknown job code"; } }</pre> <p>If the following SQLJ function is created to execute Java method <code>Jobcode.job2</code>, explicitly describing the <code>Jobcode.job2</code> signature as <code>()</code>:</p> <pre>create function job_title(jc integer) returns varchar(20) language java parameter style java external name 'Jobcode.job2()'</pre> <p>This function definition raises the following error:</p> <pre>Server Message: Number 14200, Severity 16 Server 'rosterdb', Procedure 'job_title', Line 4: The specified java signature is invalid.</pre> <p>This is because the SQLJ function signature <code>integer</code> does not map to the explicitly defined Java method signature <code>()</code>.</p>
Action	Correct either the SQLJ function signature or the Java method signature so that they can be mapped to each other.
Version in which this error is raised	Version 12.5 and later

Query Processor Errors (300s, 400s, 500s)

This chapter contains errors that result from query processing. The query processor executes specified queries. The processor yields highly efficient query plans that execute using minimal resources and ensure that results are consistent and correct. See the *Performance and Tuning Series: Query Processing and Abstract Plans* for more information about query processing.

Error	Page
Error 311	47
Error 313	49
Error 314	51
Error 403	53
Error 414	54
Error 428	55
Error 511	56
Error 512	57
Error 515	58
Error 530	60
Error 539	61
Error 540	61
Error 546	63
Error 547	64
Error 551	65
Error 584	66

Error 311

Severity

16

Message text

The optimizer could not find a unique index which it could use to scan table '%.*s' for cursor '%.*s'.

Explanation For an explicit updated cursor scan, Adaptive Server requires that a unique index exist on the table. A unique index ensures that the cursor is positioned at the correct row the next time a `fetch` is performed on that cursor. Error 311 occurs when a unique index does not exist during a scan required for a cursor marked for update.

Action Create a unique index using the following methods:

- Drop your existing index. Create an index with the `unique` option or use `alter table add constraint with the unique option`.
- Create a new index using `create index with the unique option` or use `alter table add constraint with the unique option`.

Note When you create an index with the `unique` option, if there is a duplicate key value or if more than one row contains a null value, the command is aborted and error 1505 is raised, showing the duplicate value. See to “Error 1505” on page 217 if that happens.

- Add an `IDENTITY` column using `alter table` and the `identity` option.

If your table has one or more `IDENTITY` columns, you can, for future indexes you create for this table, use the database option `identity in nonunique index`. When this option is set, any nonunique index created on a table with an `IDENTITY` column will have the `IDENTITY` column automatically included as the last key field of the index. This allows an otherwise nonunique index to be used for a cursor marked for update.

However, all indexes would then be considered unique by the optimizer since every index that is created would be unique and this could result in poor performance for some queries. Since the `identity in nonunique index` option does not affect existing indexes, only future indexes you create, the existing 311 error will not be resolved.

Warning! Do not use the `identity in nonunique index` option unless you plan to test your queries.

To use the `identity in nonunique index` option, the table must already have an `IDENTITY` column, either from a `create table` statement or by setting the `auto identity` database option to `true` before creating the table.

The commands to set this option are:

```
1> use master
2> go
```

	<pre> 1> sp_dboption <database_name>, 2> "identity in nonunique index", true 3> go 1> use <database_name> 2> go 1> checkpoint 2> go </pre>
Additional information	Refer to the <i>Reference Manual: Commands</i> for information about create index and alter table, and <i>Reference Manual: Procedures</i> for information about sp_dboption.
Version in which this error is raised	All versions

Error 313

Severity	16
Message text	A dirty read scan requires a unique index (unless the user forced a non-unique index or table scan). The optimizer could not find a unique index to use.
Explanation	Isolation level 0 (zero) allows transactions to read uncommitted data (dirty reads). Since level 0 scans do not acquire locks, modifications by other processes can cause rows being scanned to move. When this occurs, Adaptive Server restarts the scan. To restart properly, the scan must use the same key used to find the modified row, and start again from the next key in the index; that is, the scan must be on a unique index. Error 313 occurs when a unique index does not exist during a dirty read scan.
Action	<p>To resolve the problem, create a unique index using one of these methods:</p> <ul style="list-style-type: none"> • Drop your existing index. Create an index with the unique option or use alter table add constraint with the unique option.

- Create a new index using `create index` with the `unique` option or use `alter table add constraint` with the `unique` option.

Note When you create an index with the `unique` option, if there is a duplicate key value or if more than one row contains a null value, the command is aborted and error 1505 is raised, showing the duplicate value. See “Error 1505” on page 217 if that happens.

- Add an `IDENTITY` column using `alter table` and the `identity` option.

If your table has one or more `IDENTITY` columns, you can, for future indexes you create for this table, use the database option `identity in nonunique index`. When this option is set, any nonunique index created on a table with an `IDENTITY` column will automatically have the `IDENTITY` column included as the last key field of the index. This allows an otherwise nonunique index to be used for a cursor declared for update.

However, all indexes would then be considered unique by the optimizer since every index that is created would be unique and this could result in poor performance for some queries. Since the `identity in nonunique index` option does not affect existing indexes, only future indexes you create, the existing 313 error will not be resolved.

Warning! Do not use the `identity in nonunique index` option unless you plan to test your queries.

To use the `identity in nonunique index` option, the table must already have an `IDENTITY` column, either from a `create table` statement or by setting the `auto identity` database option to `true` before creating the table.

The commands to set this option are:

```
1> use master
2> go

1> sp_dboption <database_name>,
2> "identity in nonunique index", true
3> go

1> use <database_name>
2> go

1> checkpoint
2> go
```

Where `<database_name>` is the name of the actual Adaptive Server.

forceindex

Using `forceindex` with a clustered nonunique index uses key values to restart the scan at the first row that has key values.

To force Adaptive Server to use a nonunique index for your table scan, use `forceindex`. However, the restart will either be approximate or fail altogether if a row becomes invalidated. If the restart fails, your query will abort.

Warning! Do not use `forceindex` with dirty reads.

Additional information

Refer to the *Reference Manual: Commands* for information about `create index` and `alter table`, and *Reference Manual: Procedures* for information about `sp_dboption`.

Version in which this error is raised

All versions

Error 314

Severity

16

Message text

```
WARNING: A non-unique clustered index has been forced
on an isolation level 0 scan on table '%.*s'. If the
scan must restart, the scan will be repositioned at the
beginning of the duplicate key group. Thus, it is
possible that this can infinitely loop.
```

Explanation

Isolation level 0 (zero) allows transactions to read uncommitted data (dirty reads). Since level 0 scans do not acquire locks, modifications by other processes can cause rows being scanned to move. When this occurs, Adaptive Server restarts its scan. To properly restart, the scan must use the same key used to find the modified row and start again from the next key in the index. Thus, the scan must be on a unique index.

Error 314 occurs when a unique index does not exist during a dirty read scan and you have used `forceindex` to force Adaptive Server to use a nonunique clustered index for your table scan. In this case, Adaptive Server uses key values to restart the scan at the first row that has key values. However, the restart will either be approximate or fail altogether if a row becomes invalidated. If the restart fails, your query will abort. In addition, it is possible that the scan can go into an infinite loop.

Action

This is a warning message and no action is required. However, since the consequences are severe, consider using the following options instead of `forceindex`.

Create a unique index using one of the following methods:

- Drop your existing index. Create an index with the `unique` option or use `alter table add constraint` with the `unique` option.
- Create a new index using `create index` with the `unique` option or use `alter table add constraint` with the `unique` option.

Note When you create an index with the `unique` option, if there is a duplicate key value or if more than one row contains a null value, the command is aborted and error 1505 is raised, showing the duplicate value. Refer to “Error 1505” on page 217 if that happens.

- Add an `IDENTITY` column using `alter table` and the `identity` option.

If your table has one or more `IDENTITY` columns, you can, for future indexes you create for this table, use the database option `identity in nonunique index`. When this option is set, any nonunique index created on a table with an `IDENTITY` column will have the `IDENTITY` column automatically included as the last key field of the index. This allows an otherwise nonunique index to be used for a cursor declared for update.

However, all indexes would then be considered unique by the optimizer since every index that is created would be unique and this could result in poor performance for some queries. The `identity in nonunique index` option does not affect existing indexes, only future indexes you create.

Warning! Do not use the `identity in nonunique index` option unless you plan to test your queries.

To use the identity in nonunique index option, the table must already have an IDENTITY column, either from a create table statement or by setting the auto identity database option to true before creating the table.

The commands to set this option are:

```
1> use master
2> go

1> sp_dboption <database_name>,
2> "identity in nonunique index", true
3> go

1> use <database_name>
2> go

1> checkpoint
2> go
```

Additional information	Refer to the <i>Reference Manual: Commands</i> for information about create index and alter table, and <i>Reference Manual: Procedures</i> for information about sp_dboption.
Version in which this error is raised	All versions

Error 403

Severity	16
Message text	Invalid operator for datatype op: %s type: %s.
Explanation	This error occurs during Adaptive Server expression processing, when an operator in the query is applied to a datatype for which that operator is not valid. An operator is a logical or arithmetic expression such as “+” or “-”. For example: <pre>1> select pub_id * pub_name from publishers 2> go Msg 403, Level 16, State 2: Line 1: Invalid operator for datatype op: MULTIPLY type: CHAR.</pre>
Action	Correct your query and run it again.

Additional information	Refer to the Adaptive Server Enterprise reference manuals and the <i>Transact-SQL User's Guide</i> for information about the use of operators in queries.
Version in which this error is raised	All versions

Error 414

Severity	16
Message text	The current query would generate a key size of %d for a work table. This exceeds the maximum allowable limit of %d.

Explanation

When you execute a query, Adaptive Server may need to create work tables to temporarily store query results. For example, a work table is used when duplicate rows must be removed in processing a query with an aggregate function. A `sysindexes` row is built for the work table after checking that the specified command does not violate any limitations on keys for user tables.

Error 414 is raised when you execute a query containing an aggregate function, and the total length of columns named in the `group by` clause of the query exceeds the maximum limit. The maximum limit depends on the Adaptive Server page size:

Adaptive Server page	Maximum bytes
2K	600
4K	1250
8K	2600
16K	5300

Action	<p>Check the command for possible violations of the maximum size limit. You can correct the problem by doing one of the following:</p> <ul style="list-style-type: none"> • Reduce the columns named in the <code>group by</code> clause until the combined column length is no more than the maximum number of bytes for the Adaptive Server page size. • If it is necessary to group by a large <code>character</code> column, consider using the <code>substring</code> string function on the column in the <code>group by</code> clause. This allows Adaptive Server to build a composite key to group the result set using only a portion of the character string. For example:
---------------	---

```
1> select * from titles
2> group by title, substring(notes,1,10)
```

This query uses only the first 10 bytes of *notes* (a 200 byte varchar column) to group the data.

Version in which this error is raised

All versions

Error 428

Severity

20

Message text

There are more than %d referential constraints on table %.*s. Please reduce the number of referential constraints before trying this query.

Explanation

During query processing, Adaptive Server checks for the existence of foreign keys and dependent foreign keys (a foreign key is a column or combination of columns whose values match the primary key). A range table entry is created for each reference check and foreign key constraint. “%d” in the error message is the maximum number of table references allowed. Error 428 is raised when this limit is reached.

Action

When setting up constraints on your tables, determine the maximum number of tables that might be touched by an update, insert, or delete statement.

To determine which constraints exist for a table, type:

```
1> use <database_name>
2> go
```

```
1> sp_helpconstraint <table_name>
2> go
```

If necessary, drop some of the constraints on the table:

```
1> alter table <table_name>
2> drop constraint <constraint_name>
3> go
```

Additional information

Refer to the Adaptive Server Enterprise reference manuals for information about constraints.

Version in which this error is raised

All versions

Error 511

Severity 16

Message text **Version 12.5 and later**

Attempt to update or insert row failed because resultant row of size %d bytes is larger than the maximum size (%d bytes) allowed for this table.

Version 12.0.x and earlier

Updated or inserted row is bigger than maximum size (%d bytes) allowed for this table.

Explanation

This error occurs when you try to insert or update a row that is longer than the maximum allowable length. The maximum allowable row size that includes row overhead depends on the Adaptive Server page size:

Adaptive Server page max	APL row length max	DOL row length
2K	1962	1964
4K	4010	4012
8K	8016	8108
16K	16298	16298

The use of variable or fixed columns will alter the size of the overhead required for a row.

Error 511 is caused by database design errors (for example, a table designed with the potential for rows containing more than the maximum allowable characters). The following warning is given when you create a table that has the potential for rows exceeding the maximum row size (that is, the maximum length of all columns added up is greater than the allowable number of characters):

Msg 1708, Level 16, State 1:Warning: Row size could exceed row size limit, which is %d bytes.

Action

Use one of these actions to resolve the error:

- If error 511 is caused by a table containing rows with more than the maximum row size, divide the table into two or more tables so that no row length is more than the allowable number of characters.
- If error 511 error does not appear to be caused by the above situation, call Sybase Technical Support.

Version in which this error is raised All versions

Error 512

Severity 16

Message text Subquery returned more than 1 value. This is illegal when the subquery follows =, !=, <, <=, >, >=, or when the subquery is used as an expression.

Explanation When an expression subquery returns more than one result, it violates the relational operator rule for the outer query, and error 512 occurs. An example of an expression subquery that returns one result follows:

```
1> select * from table_one where x =
2> (select sum(a) from table_two
3> where b = table_one.y)
4> go
```

An example of a query that returns more than one result and causes error 512 follows:

```
1> use pubs2
2> go

1> select authors.au_id from authors where
2> authors.au_id = (select titleauthor.au_id
3> from titleauthor)
4> go
```

```
Msg 512, Level 16, State 1:
Line 1:
Subquery returned more than 1 value. This is illegal
when the subquery follows =, !=, <, <=, >, >=, or when
the subquery is used as an expression.
```

Action To correct the problem in the example, use “in” in place of “=”, as in the following example:

```
1> select authors.au_id from authors where
2> authors.au_id in (select titleauthor.au_id
3> from titleauthor)
4> go
```

```

au_id
-----
172-32-1176
213-46-8915
.
.
899-46-2035
998-72-3567

```

(19 rows affected)

Additional information	Refer to the chapter “Subqueries: Using Queries Within Other Queries” in the <i>Transact-SQL User’s Guide</i> .
Version in which this error is raised	All versions

Error 515

Severity	16
Message text	Attempt to insert NULL value into column '%.*s', table '%.*s'; column does not allow nulls. Update fails.%S_EED

Explanation When you create a table, you can explicitly define whether each column should allow null values. If you do not specify NULL or NOT NULL for a column when you create the table, the default value will be NOT NULL. If you use `sp_dboption` to set allow nulls by default to TRUE for the database, any new table that is created will have columns with the default value NULL.

Error 515 occurs at run time when a column has a NOT NULL default and you try to insert a NULL value into that column (for example, if a user does not enter a value for that column). The error message includes:

- The name of the affected column.
- The name of the affected table.
- The extended error data array (EED), which includes the database name, owner, table name, and column name. This information does not display, but you will see it in the message definition if you select error 515 from `sysmessages`.

The following example would cause a 515 error:

```
1> create table table1 (column1 varchar (5))
```

```

2> go

1> declare @c varchar(5)
2> insert into table1 values (@c)
3> go

```

Msg 515, Level 16, State 3:
 Server 'SERVER_NAME', Line 2:
 Attempt to insert NULL value into column 'column1', table 'test.dbo.table1';
 column does not allow nulls. Update fails.

Action To determine whether a column has NULL or NOT NULL defined, enter commands like the following, replacing “pubs2” with the name of the database in which your table resides and “titleauthor” with the table in which your column resides:

```

1> use pubs2
2> go

1> sp_help titleauthor
2> go

```

Name	Owner	Type						
titleauthor	dbo	user table						
Data_located_on_segment		When_created						
default		Oct 27 2007 10:09AM						
Column_name	Type	Length	Prec	Scale	Nulls	Default_name	Rule_name	Identity
au_id	id	11	NULL	NULL	0	NULL	NULL	0
title_id	tid	6	NULL	NULL	0	NULL	NULL	0
au_ord	tinyint	1	NULL	NULL	1	NULL	NULL	0
royaltyper	int	4	NULL	NULL	1	NULL	NULL	0

The Nulls column indicates whether null values are allowed. A value of 0 (zero) for the column means nulls are not allowed and 1 means null values are allowed.

To change the default for a column from NOT NULL to NULL, use the alter table command to modify the column. For example:

```
alter table table1 modify column1 varchar(5) null
```

If you want new tables in the database to have the default NULL, use the following commands for the database:

```

1> use master
2> go

```

```
1> sp_dboption <database_name>,  
2> "allow nulls by default", true  
3> go  
  
1> use <database_name>  
2> go  
  
1> checkpoint  
2> go
```

Where *<database_name>* is the name of the database with the behavior you want to change.

Version in which this error is raised

All versions

Error 530

Severity

16

Message text

```
Attempt to insert NULL value into column %d in work  
table (table id %ld); column does not allow NULLS.  
UPDATE fails.
```

Explanation

During the runtime phase of updates, Adaptive Server sometimes uses worktables to temporarily store query results. Later in the query processing, Adaptive Server selects the values from those tables.

If, as a result of your query, Adaptive Server tries to insert a null value into a column of a worktable, and nulls are not allowed for the column, error 530 occurs. The command is aborted and the update fails. Error 530 is caused by an Adaptive Server problem.

Action

Depending on the context of the error, you may be able to rewrite the query (for example, supply a column name or specify `isnull`). If the error reoccurs, or you cannot write the query in another way, call Sybase Technical Support.

Version in which this error is raised

All versions

Error 539

Severity	20
Message text	Unexpected internal access methods error %d, state %d. Please report to Technical Support the following information: dbid=%d, objectid=%ld, curcmd=%d (%s), plasteerror=%ld, pstat=0x%x, p2stat=0x%x, xactid=(%ld, %d) .
Explanation	Error 539 is raised during query processing when Adaptive Server attempts to execute an access method, but the attempt fails with an unknown failure in the access method. This is why error 539 reports error code 0. Error 539 is due to an Adaptive Server problem.
Action	If additional errors accompany the 539 error, resolve them using the relevant write-ups in this document. If the problem persists, call Sybase Technical Support.
Additional information	Have the information listed in “Reporting errors” on page 4 when you call Sybase Technical Support:
Version in which this error is raised	All versions

Error 540

Severity	16
Message text	Schema for table '%.*s' has changed since compilation of this query. Please re-execute query.
Explanation	<p>A table’s schema consists of column definitions along with any indices, constraints, rules, defaults and so on. A change to any of these elements (for example, creating or dropping a constraint) constitutes a change in the table’s schema.</p> <p>Query processing involves these basic steps:</p> <ul style="list-style-type: none"> • Adaptive Server examines the SQL statement and, if it passes syntax checks, the query is parsed to produce an internal representation called the query tree.

- The server then determines the optimal access strategy to implement the query and compiles the query tree to produce the query plan. For ad hoc queries the server discards the query tree. For stored procedures, the query tree is saved on disk so that the plan can be regenerated as necessary.
- The query plan is executed to produce the results.

Error 540 is raised:

- When the server attempts to execute the query plan of an ad hoc query, but discovers that the table's schema changed after the query tree was produced and before the plan has been executed. Adaptive Server cannot recompile an ad hoc query because its tree has not been preserved. On the other hand, a stored procedure will always be recompiled after the server detects a schema change. Note that schema changes to a table include updating the table statistics.
- When you create a table with a primary key constraint on a column, followed by an insert into that table in the same batch. For example:

```
1> create table table1 (id int primary key, f2 int )
2> insert into table1 (id, f2) values (1,1)
3> go
```

```
Msg 540, Level 16, State 1: Server 'tsg125', Line 2:
Schema for table 't1' has changed since compilation
of this query. Please re-execute query.
```

In this scenario, error 540 occurs due to an Adaptive Server problem.

Error 540 is not a serious error; however, it may prevent you from running some ad hoc queries.

Action

If the error occurred when creating and inserting to a table with a primary key constraint in the same batch, contact Sybase Technical Support to help you upgrade to a version in which the problem is resolved.

If the error occurred in a different scenario:

- Retry the query. This ensures that Adaptive Server can build the appropriate query plan for the revised schema.
- Avoid taking actions which result in schema changes to tables (for example, running `update statistics`) while queries referencing those tables are executing.

Version in which this error is raised

All versions

Error 546

Severity	16
Message text	Foreign key constraint violation occurred, dbname = '%.*s', table name = '%.*s', constraint name = '%.*s'.%S_EED
Explanation	Foreign key constraints are a form of integrity constraint which ensure that no insert or update on a foreign key table is performed without a matching value in the primary key table. Error 546 is raised when an integrity constraint is violated during execution of a query. For example, error 546 can be raised when Adaptive Server detects that the data inserted into a foreign key does not match any primary key value in the table referenced by the foreign key.

For example:

```
1> create table departments
2>   (d_id int primary key, d_name varchar(30))
3> create table employees
4>   (emp_id int, empname varchar(100), d_id int references departments)
5> go
```

```
1> insert departments values (1, 'sales')
2> go
```

(1 row affected)

```
1> insert employees values (1, 'Fred Smith', 1)
2> go
```

(1 row affected)

```
1> update employees set d_id = 2 where emp_id=1
2> go
```

```
Foreign key constraint violation occurred, dbname = 'master', table name =
'employees', constraint name = 'employees_d_id_1824006498'.
```

Command has been aborted.

(0 rows affected)

In this example, error 546 is raised because the update command attempted to use a non-existent department.

Action Check your query or application to determine the source of the constraint violation. You can examine the constraint by executing:

```
1> sp_helpconstraint <table_name>
```

```
2> go
```

Where `<table_name>` is the table appearing in the 546 message.

Version in which this error is raised

All versions.

Error 547

Severity

16

Message text

```
Dependent foreign key constraint violation in a referential integrity constraint. dbname = '%.*s', table name = '%.*s', constraint name = '%.*s'.%S_EED
```

Explanation

Adaptive Server provides integrity constraints to help you maintain logical data integrity in a database. Referential integrity (or foreign key) constraints are a type of constraint which require that data being inserted into a given table column already has matching data in another column (the target column), which may be in the same table or another table. The column on which the constraint is declared can be considered the child, or dependent, in a parent-child relationship.

Error 547 is raised when updating or deleting rows from a parent table would remove target columns matching dependent data in the child table. For example:

```
1> create table parent
2> (a int primary key,
3> b int,
4> unique (b))
5> go

1> create table child
2> (c int primary key,
3> d int references parent(b))
4> go

1> insert parent values (11,22)
2> insert parent values (13,26)
3> insert child values (101,22)
4> go

1> delete parent where a=13
2> go
```

```
(1 row affected)
1> delete parent where a=11
2> go
```

Msg 547, Level 16, State 1: Line 1:
 Dependent foreign key constraint violation in a referential integrity constraint. dbname = 'hrdb', table name = 'parent', constraint name = 'detail_d_1088006907'.
 Command has been aborted.
 (0 rows affected)

```
1> update parent
2> set b=29
3> where a=11
4> go
```

Msg 547, Level 16, State 1: Line 1:
 Dependent foreign key constraint violation in a referential integrity constraint. dbname = 'hrdb', table name = 'parent', constraint name = 'child_d_1088006907'.
 Command has been aborted.
 (0 rows affected)

Notice that you may drop or update rows in parent provided you do not affect the referential integrity constraints.

Action	Delete or update the dependent data in the child table before deleting or changing the parent data.
Additional information	Use the system procedure <code>sp_helpconstraint</code> to view the referential constraints in affect for a table. See the <i>Reference Manual: Procedures</i> for detail about using the <code>sp_helpconstraint</code> procedure.
Version in which this error is raised	All versions

Error 551

Severity	20
Message text	An unknown EVAL was sent to the execution module.
Explanation	The instructions Adaptive Server creates to run a query are in an evaluation list that contains instruction-argument pairs. Error 551 occurs when Adaptive Server receives an illegal instruction. It is probably caused by an incorrectly compiled query tree.

Action	<p>Drop and re-create the procedure or trigger being run:</p> <ol style="list-style-type: none">1 If you do not have a script for re-creating your procedure or trigger, get the text of the procedure or trigger:<pre>1> use <database_name> 2> go 1> sp_helptext <object_name> 2> go</pre><p>Where <i><database_name></i> is the name of the database in which the procedure or trigger resides and <i><object_name></i> is the name of the procedure or trigger.</p><p>Alternatively, use the <i>defncopy</i> program to copy the procedure or trigger definition to a file. Refer to “<i>defncopy</i>” in the chapter “Utility Commands Reference” of the Adaptive Server Enterprise 15.0 <i>Utility Guide</i> for details.</p>2 Drop the procedure or trigger:<pre>1> drop procedure <object_name> 2> go</pre><p>or:</p><pre>1> drop trigger <object_name> 2> go</pre>3 Re-create the procedure or trigger. Refer to “create procedure” and “create trigger” in the <i>Reference Manual: Commands</i> for instructions.
Additional information	If error 551 persists, create a scenario to reproduce the problem and contact Sybase Technical Support. Be sure to have the information on hand listed in “Reporting errors” on page 4, including a copy of the query that raised the error.
Version in which this error is raised	All versions

Error 584

Severity	16
Message text	Explicit value specified for identity field in table '%.*s' when IDENTITY_INSERT is set to OFF.

Explanation	<p>Each table can include one <code>IDENTITY</code> column. <code>IDENTITY</code> columns store sequential numbers that are generated automatically by Adaptive Server. The value of the identity column can uniquely identify each row in a table.</p> <p>The Adaptive Server query processing option <code>set identity_insert</code> determines whether explicit inserts into a table's <code>IDENTITY</code> column are allowed. Inserting a value into the <code>IDENTITY</code> column allows you to specify a "seed" value for the column or to restore a row that was deleted in error. Setting <code>identity_insert on</code> allows the table owner, database owner, or System Administrator to explicitly insert a value into an <code>IDENTITY</code> column. Unless a unique index exists on the <code>IDENTITY</code> column, any positive value without regard to uniqueness may be inserted in <code>IDENTITY</code> columns when <code>identity_insert</code> is set to <code>on</code>.</p> <p>Setting <code>identity_insert off</code> prohibits inserts to <code>IDENTITY</code> columns.</p> <p>Error 584 is raised if you attempt to insert an explicit value into an <code>IDENTITY</code> column when <code>identity_insert</code> is set to <code>off</code>.</p>
Action	<p>Set <code>identity_insert on</code> before attempting to insert an explicit value into an <code>IDENTITY</code> column. <code>identity_insert</code> should be reset to <code>off</code> once the insert operation is complete.</p>
Additional information	<p>The syntax for setting the <code>identity_insert</code> option is:</p> <pre>1> set identity_insert '<table_name>' {on off} 2> go</pre> <p>Where <code><table_name></code> is the base table for the column. Only the table owner, database owner, or System Administrator can set this option.</p>
Version in which this error is raised	<p>All versions</p>

Access methods form the interface between the server and an external object. This section contains error messages that result from various problems when Adaptive Server uses various access methods.

Error	Page
Error 601	69
Error 603	71
Error 605	72
Error 611	75
Error 614	76
Error 622	80
Error 623	81
Error 624	83
Error 625	86
Error 629	89
Error 631	91
Error 644	95
Error 678	96
Error 691	97
Error 692	98
Error 693	100
Error 694	101
Error 695	103
Error 696	105
Error 697	106

Error 601

Severity

21

Message text

Descriptor for system table '%ld' in database '%d'
not found in the descriptor hash table.

Explanation	Adaptive Server expects every database to have a full set of system tables. Error 601 occurs when the Adaptive Server Threshold Manager could not open a system table in a database because the descriptor for the database was not found. The failure to open the system table is usually due to corruption in the database named in the error message.
Action	<ol style="list-style-type: none">Determine the name of the database from the database ID in the error message:<pre>1> use master 2> go 1> select * from sysdatabases where dbid = <ID> 2> go</pre>Where <i><ID></i> is the database ID from the error message.Determine whether the system table object ID displayed in the error message exists:<pre>1> use <database_name> 2> go 1> select * from sysobjects where id = <object_ID> 2> go</pre>Where <i><database_name></i> is the name from step 1 and <i><object_ID></i> is the system table ID listed in the error message.If the ID exists for the table displayed in the message, then error 601 may be caused by corruption in the memory structure used to hold the descriptor. Shutting down and restarting Adaptive Server should clear the problem. If this does not work, go to step 4.If the ID does not exist, there is probably corruption in your database. Do the following:<ol style="list-style-type: none">Run <code>dbcc checkalloc</code>, <code>dbcc checkcatalog</code>, and <code>dbcc checkdb</code> for the database.Determine whether hardware problems exist by checking your operating system error log.Call Sybase Technical Support to help you recover your database, but you will probably have to recover your database from backups.
Additional information	Before contacting Technical Support, have on hand the information listed in “Reporting errors” on page 4, including the output from <code>dbcc checkalloc</code> , <code>dbcc checkcatalog</code> , and <code>dbcc checkdb</code> for the database listed in the error message.

Version in which this error is raised All versions

Error 603

Severity 17

Message text There are not enough system session descriptors available to run this query. Raise the value of the configuration parameter 'number of aux scan descriptors' or retry your query later.

Explanation This error occurs when Adaptive Server runs out of system session descriptors when running a query. Error 603 can occur on commands which create, drop, and update objects, as these commands often have many dependencies on system tables which Adaptive Server must handle to complete the command.

Session descriptors are the data structures used to manage access to Adaptive Server objects. Session descriptors come in three forms—user session descriptors (including worktable session descriptors), system session descriptors, and worktable session descriptors.

Session descriptors are used when:

- A query is compiled or executed, a user session descriptor is reserved and assigned to each table involved in the query. Adaptive Server allocates user session descriptors so that the number of tables involved in a query is limited to 50 user tables and 14 work tables. If a single table is referenced more than once in the *from* clause (a “self join”), a user session descriptor is reserved for each reference. User session descriptors correspond to the number of open objects configuration parameter.
- System session descriptors are used for system tables that are not named in the query. For example, the following query prompts Adaptive Server to open `sysobjects` and `syscolumns` to look up the table `t` and its columns `x` and `y`:

```
select x from t where y = 3
```

A maximum of six system session descriptors are available.

- A worktable session descriptor is used to manage references to worktables. A worktable is required for order by, group by, distinct, and some other operators. Also, the query optimizer may choose a plan involving reformatting which copies rows from a table into a worktable, creates an index, then performs joins between the worktable and other tables in the query. There are 12 worktable session descriptors that can be used at any one time.

You cannot change the number of session descriptors available to a query.

Action

When you error 603 is raised, you must simplify your query or command. To recover from this error, split up your query or simplify the command and run it again.

For example, if you are entering a drop procedure command and that procedure contains other procedures or triggers, you can simplify your command by dropping the other procedures or triggers manually before running the drop procedure command.

Additional information

Refer to the chapter “Setting Configuration Parameters” in the *System Administration Guide: Volume 1* for a description of “number of aux scan descriptors” and how queries use scan descriptors.

Version in which this error is raised

All versions

Error 605

Severity

21

Message text

Versions earlier than 15.0:

```
An attempt was made to fetch logical page '%ld' in
database '%.*s' from cache '%.*s'. Page belongs to
object with id '%ld', not to object '%.*s'.
```

Version 15.0 and later:

```
An attempt was made to fetch logical page '%ld' from
cache '%.*s'. Page belongs to %S_PTININFO and not to
%S_PTININFO.
```

Explanation This serious error occurs when Adaptive Server discovers page allocation corruption. Adaptive Server tries to access a particular object but discovers a page in the page chain for that object whose object ID is different than that of the object being accessed. There is probably either a damaged page chain or an invalid entry in the system tables for that object.

Usually this error is detected after the corruption has been written to the database on disk (hard 605), but it can also occur entirely in cache without the damage ever being written to disk (transient 605), and is not associated with data corruption. See the next section for more information about the different kinds of 605 errors and their causes.

Note When an error is transient, it means it is reported even though no error condition exists or it exists only in memory. Generally, a transient error is cleared by a reboot of Adaptive Server, whereas a hard error does not go away with an Adaptive Server reboot.

For accurate diagnosis of object ID mismatch problems, the 605 error is raised only when the objects in question are user objects. Other related types of errors result in one of the errors 691 through 697. For details on these errors, refer to the write-ups in this chapter.

How much corruption exists?

Assuming the error is not a transient 605, corruption usually exists in the second object specified in the 605 error text. If the `dbcc checkdb` and `dbcc checkalloc` commands do not report additional errors, the first object mentioned is not corrupt. Because the 605 error can mask the existence of other errors, you must run the `dbcc` checks to determine the extent of the damage.

Instead of listing two object names, this error might contain other information:

- The error displays a number greater than zero. An attempt was made to refer to an object ID that does not exist in the system table `sysobjects`.
- If the error states that a page belongs to object “ALLOCATION,” some of the allocation structures used by the database may be corrupted.

The following table lists hardware causes for 605 errors.

Hardware cause	Additional information	Transient or hard
Overlapping partitions on disk (UNIX only).	Often this occurs when the server is installed on a partition that overlaps other partitions.	Hard
Note See the <i>Configuration Guide for UNIX</i> for more information about partitions.		

Hardware cause	Additional information	Transient or hard
Problem with device driver	--	Hard
Problem with controller	--	Hard
Bad disk	“Data retry” messages in <i>/var/adm/messages</i> (UNIX).	Hard

The following table lists software causes for 605 errors.

Software causes	Additional information	Transient or hard
Loading a master database on a master device that does not have rows in <i>sysusages</i> for <i>dbid</i> 1,2,3 identical to the device it was dumped from.	--	Hard
Loading a master database that has <i>syscharsets</i> in a location that is different from the default.	Caused by changing Adaptive Server sort order after operating it for some time.	Hard
Operating system notifies Adaptive Server that an I/O operation is complete when it is not. Any time buffer information about a page is different from the disk information about the page (and the disk has the correct information).	Adaptive Server connection often dies. Running <i>dbcc checkdb</i> and <i>dbcc checkalloc</i> should produce no more 605 errors. Corruption is not on disk. First, <i>checkpoint</i> the other databases. Then, perform a shutdown...with <i>nowait</i> and restart Adaptive Server.	Transient

Action

To resolve the error:

- 1 Run the *dbcc checktable* command on the second object specified in the error message.
- 2 To help determine the full extent of the corruption, run the *dbcc checkdb* and *dbcc checkalloc* commands as soon as possible.
- 3 Check the Adaptive Server error log for other errors that often accompany a 605 error. If the 605 error is not transient, the problem is severe and you probably need to restore from known clean backups.
- 4 Either of the following situations may also apply to your Adaptive Server:
 - If this error occurs after rebuilding the master device or reconfiguring the configuration block to the defaults, then the sort order on the new device may not match the sort order in the dump. If the sort orders do not match, call Sybase Technical Support for further assistance.
 - To check if the error occurred as the result of hardware failure, examine your operating system error log and correct hardware problems.

Additional information For information on locating device fragments, use the procedures in Determining the physical device on which a database resides in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.

Have the information listed in “Reporting errors” on page 4 on hand before contacting Technical Support, including the output of dbcc checkdb, dbcc checkalloc, and dbcc page for the corrupted page.

Related page mismatch errors

Starting with version 11.0.3, many page mismatch problems no longer raise 605 errors, but are reported with the following error numbers:

Error reference	Description
“Error 691” on page 97	Page 0 is read incorrectly, or an attempt is made to read a negative page number.
“Error 692” on page 98	Attempt to read an uninitialized page.
“Error 693” on page 100	I/O request was made prematurely, causing an object ID mismatch.
“Error 694” on page 101	The device or operating system may be suspect.
“Error 695” on page 103	The page has been overwritten, for example page 10 on disk may actually hold page 20. Hard error.
“Error 696” on page 105	Objects are crossed in tempdb. Similar to a 605 except that objects are in tempdb.
“Error 697” on page 106	Results from memory corruption.

Refer to the write-ups for these errors in this manual for more information.

Version in which this error is raised All versions

Error 611

Severity 21

Message text Attempt made to end a transaction that is idle or in the middle of an update.

Explanation This error occurs when a termination request event occurs during the processing of a transaction. Error 611 can be caused by:

- Massive processing using group by or order by
- A stored procedure that references other stored procedures (and is in the midst of calling those procedures) being dropped

Action	<ul style="list-style-type: none">• An Adaptive Server problem To determine what caused the error, check for: <ul style="list-style-type: none">• <code>dbcancel</code> calls from within DB-Library exit processing routines that fail to test all error conditions• Users typing <code>Ctrl-C</code> during heavy use of temporary tables Error 611 can produce a stack trace that shows some of the SQL that caused the error. Obtain the entire stack trace and any errors that occurred prior to or after the 611 error, to determine the internal process state at the time of the failure.
Version in which this error is raised	All versions

Error 614

Severity	21
Message text	<p>Versions earlier than 15.0:</p> <pre>A row on page %ld was accessed that has an illegal length of %d in database '%.*s'.</pre> <p>Version 15.0 and later:</p> <pre>Adaptive Server accessed a row that has an illegal length of %d while in database '%.*s'. %S_PAGE. The minimum row length is %d. The page size is %d.</pre>

Note This error may be caused by a hardware problem.

Explanation	<p>This error occurs when Adaptive Server accesses a data or index row whose length is smaller than the minimum row size or greater than the maximum row size.</p> <p>The minimum length of a row for each object is stored in the <code>minlen</code> column of <code>sysindexes</code> and in each data or index page header. The size of a data row or index row is 1962 characters; allow 2 characters of row overhead for all-pages-locked (APL) tables, for a maximum usable row size of 1960. For data-only-locked (DOL) tables, subtract two characters per varchar column in determining usable row size.</p> <p>This error can occur:</p>
-------------	---

- During normal processing, when Adaptive Server tries to access the row specified by the error message
- During database recovery (database recovery occurs during Adaptive Server start-up or a load database or load transaction command is processed)

Error 614 can be caused by data corruption during normal processing (for example, an operating system panic occurs, causing interruption in disk writes when using UNIX files for Sybase database devices). This may be due to a problem with Adaptive Server, the operating system, or hardware.

Action

Error 614 is usually the result of a more serious underlying problem, and recovering from this error depends on when the error occurred. Determine whether the error occurred during normal processing or during database recovery, then follow the appropriate set of instructions in this section.

If the error occurred during normal processing

- 1 Use the procedure in “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* to identify the table and index that correspond to the page number from the error message text.

Warning! The procedure referenced in the step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 If the object encountering the error is not a system table (a system table’s object ID is less than 100), go to step 3.

If the object encountering the error is a system table and the index ID is not 0 (zero), see “Fixing a system table corrupted index” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.

If the index ID is 0, contact Sybase Technical Support to help you repair the corruption; however, it may be necessary to restore the database from clean backups.

- 3 For user tables, if the index ID is 0, go to step 4.

- a If the index ID is not 0, translate the ID into an index name:

```
1> use <database_name>
2> go

1> select name from sysindexes
2> where id = <object_ID> and indid = <index_ID>
3> go
```
 - b To ensure that the information needed to re-create the index is available, run `sp_helpindex` on the index prior to dropping it.
 - c Drop the index.
 - d Re-create the index. This clears the corruption in most cases.
 - e Run `dbcc checktable` on the table to verify that the corruption is gone.
- 4 If the index ID is 0, perform one of the following steps:
- Restore the database from clean backups.
 - See the section “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.

Warning! Some data might be lost on this page if you recover your table using `bcp` or `select into` (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to see which rows are different (corrupted).

Before dumping a database, verify that it works correctly by running these commands prior to each dump:

```
1 dbcc checkdb.
```

- 2 dbcc checkalloc or dbcc checkalloc with the fix option. (See “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for information about how to run these commands in multi-user mode and how to prevent spurious allocation errors from dbcc commands.)

Note Alternatively, using dbcc checkstorage with checkverify can examine for most types of consistency problems in less time.

If the error occurred during database recovery

When this error occurs during recovery, the database is marked suspect and is not accessible. Usually, you must load the database from backup. To do this, follow the instructions below:

- 1 If the database with the 614 error is tempdb, check the page number. If this page is not valid for tempdb, the problem may be on another database (possibly due to a bad device). Stop here and contact Sybase Technical Support.

If the page is valid for tempdb, or the 614 error refers to a database other than tempdb, continue with step 2.

- 2 Drop the database. If the drop fails, follow the instructions in “Dropping a database when drop database fails” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.
- 3 Create a database for load. Make sure the database you create has sizes at least as large as those in sysusages for the original database (and that all other sysusages values match the original values). Refer to “create database” in the *Reference Manual: Commands* for information on “create database for load.”
- 4 Load the database from backup. (Refer to “load database” in the *Reference Manual: Commands*.)
- 5 Use the “online database” command to make the database available for use.

If loading from backups is not feasible, call Sybase Technical Support with the information from “Reporting errors” on page 4 available, including the output from `select * from master..sysusages`.

Preventing error 614 on recovery

To prevent error 614 occurring on recovery, checkpoint each database that is being used before shutting down Adaptive Server.

Additional information	Refer to “Developing a Backup and Recovery Plan” in the <i>System Administration Guide: Volume 2</i> for information about how to safely create, dump, load, and re-create databases.
Version in which this error is raised	All versions

Error 622

Severity	20
Message text	<p>Versions earlier than 15.0:</p> <pre>Opentable was passed a varno of %d. Object '%.*s' already has that session descriptor in use.</pre> <p>Version 15.0 and later</p> <pre>Opentable was passed a varno of %d. Object '%.*s' in database '%.*s' already has that session descriptor in use. This occurred while opening object id %d in database id %d.</pre>
Explanation	<p>A session descriptor is an internal data structure in Adaptive Server which contains information about a table. All open objects in Adaptive Server require a descriptor structure. Whenever a connection to Adaptive Server attempts to open a table, it gets a session descriptor. <i>varno</i> refers to a variable number that Adaptive Server uses to maintain an array of session descriptors for a query. A unique <i>varno</i> is assigned to each session descriptor tracking the usage of a given object for the query.</p> <p>Error 622 is raised when Adaptive Server attempts to open a table, but the <i>varno</i> points to a session descriptor that is already in use.</p> <p>Error 622 breaks the user’s connection to Adaptive Server.</p>
Action	Error 622 is due to an Adaptive Server problem. Contact Sybase Technical Support for assistance when this error is raised.
Additional information	When calling Technical support have the information on hand listed in “Reporting errors” on page 4, including the text of the query that raised the error.
Version in which this error is raised	All versions

Error 623

Severity 21

Message text	<p>Versions prior to 15.0:</p> <pre>Attempt to retrieve row from page via RID failed because logical page %ld is not a data page. %S_RID. %S_PAGE.</pre> <p>Version 15.0 and later:</p> <pre>Adaptive Server failed to retrieve a row via its RID in database '%.*s' because the target page is not a data page. %S_RID. %S_PAGE.</pre>
	<hr/> <p>Note This error may be caused by a hardware problem.</p> <hr/>
Explanation	<p>This error occurs when Adaptive Server tries to retrieve a row from a data page by specifying the row ID (RID) and the retrieval fails because the requested page was not a data page. This error is usually caused by a server problem.</p> <p>After loading a database on top of an existing database, error 623 may appear during recovery or while doing a checkpoint.</p> <p>This error may or may not be transient (transient errors disappear after Adaptive Server is restarted). In the worst case, error 623 may indicate corruption in your database. Perform the action below as soon as possible to determine how serious the error is and to avoid further corruption if corruption has occurred.</p>
Action	<p>Determine whether this error is transient by performing these steps:</p>

- 1 Refer to “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the *Troubleshooting and Disaster Recovery* guide for instructions on how to use the `dbcc page` command to determine the table involved. Save this information and use it in the appropriate section below.

Warning! The procedure referenced in the step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 Use the object ID found in step 1 to issue this query to determine the table name:

```
1> select object_name(<object_ID>)
2> go
```

- 3 Use the table name displayed in the output from step 2 and issue this query to determine whether corruption has occurred:

```
1> dbcc checktable(<table_name>)
2> go
```

If no errors occur when you run `dbcc checktable`, the 623 error was probably transient. Shut down and restart the server.

If the error is not transient

If errors do occur when you run `dbcc checktable`, refer to the write-ups for those errors in this manual for instructions on how to recover from them or restore your database from a clean backup. To restore from backup, follow the instructions in the section “Restoring a database from backups” in the chapter “Developing a Backup and Recovery Plan in the *System Administration Guide: Volume 1*.”

If loading from backups is not feasible, call Sybase Technical Support. Have the information on hand listed in “Reporting errors” on page 4, including the output from `select * from master..sysusages`.

Version in which this error is raised

All versions

Error 624

Severity

21

Message text

Versions prior to 15.0:

```
Attempt to retrieve row from page via RID failed
because the requested RID has a higher number than the
last RID on the page. %S_RID.%S_PAGE.
```

Version 15.0 and later:

```
Adaptive Server failed to retrieve a row via its RID in
database '%.*s' because the requested RID has a higher
number than the last RID on the page. %S_RID. %S_PAGE.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server tries to retrieve a row from a data page by specifying the row ID (RID), but the retrieval fails because the requested row ID was a higher number than the last row ID on the page. This can happen:

- During normal processing, if the leaf page of a corrupt non clustered index points to an incorrect or nonexistent row ID on a data page.
- During database recovery at Adaptive Server startup (database recovery also occurs when a `load database` command is processed). Recovery fails because the data structure which contains the last checkpoint record in the transaction log is pointing to an incorrect or nonexistent record in the log.

Errors 624 and 625 occur under similar conditions.

Action

Recovery from this error depends on when the error occurred. Determine whether the error occurred during normal processing or during database recovery, then follow the instructions in one of these sections:

- “If the error occurred during normal processing” on page 84
- “Index unknown, query known” on page 84
- “Index and query both unknown” on page 85
- “if the error occurred during database recovery” on page 85

If the error occurred during normal processing

Refer to “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the *Troubleshooting and Disaster Recovery* guide for instructions about how to use the `dbcc page` command to determine the table involved. Save this information and use it in the appropriate section below.

The specific action you take depends on whether or not you know which index or query caused the problem. In general, dropping and re-creating the index should resolve the index corruption.

Index unknown, query known

If you do not know which index is causing the problem, but you do know which query encounters the problem, use this section to correct it. If you do not know the index or query, go to “Index and query both unknown” on page 85.

- 1 Determine which index should be dropped by reading the `showplan` output for the query that encounters the error and determining which nonclustered index the query is using to access the table in question.

Turn on `showplan` and use the `no execute` mode:

```
1> set showplan on
2> go

1> set noexec on
2> go
```

Enter the query that was causing the problem. For example:

```
1> select title from titles where title > "Cooking"
2> go

STEP 1
The type of query is SELECT
FROM TABLE
titles
Nested iteration
Index: titlein
```

Turn `showplan` and the `no execute` mode off again:

```
1> set noexec off
2> go

1> set showplan off
2> go
```


- 2 If the index identified in step 1 is on a system table (object ID is less than 100), refer to “Fixing a system table corrupted index” in the chapter “Other Useful Tasks” in the *Troubleshooting and Disaster Recovery* guide for instructions on repairing the system table index. Otherwise, drop and re-create the index identified in step 1 (in this example “titleind”).

Index and query both unknown

If you do not know either the index or the query, rebuild all nonclustered indexes on the table using the instructions in this section.

- 1 Examine the dbcc page output you obtained earlier.

If the object is a system table (the object ID is less than 100), refer to “Fixing a system table corrupted index” in the “Other Useful Tasks” chapter of the *Troubleshooting and Disaster Recovery* guide for instructions on how to repair the system table index. Then go to step 3.

- 2 For a user table, use `sp_helpindex` to list all indexes on the table, and then rebuild all the nonclustered indexes using one of two methods:

- Drop and re-create each nonclustered index on the table.

or:

- If a clustered index also exists on the table, drop and re-create it; this causes all nonclustered indexes to be rebuilt. If your table is large, you may not have the space to do this (a rough rule of thumb is 120 percent to 150 percent of your table size must be available).

- 3 Run `dbcc checktable` on the table to confirm that the problem has been resolved.

if the error occurred during database recovery

When this error occurs during recovery, the database is marked suspect and is not accessible. Usually, you must load the database from backup. Follow the instructions below, depending on the version level of your server:

- 1 Drop the suspect database using the instructions in “Dropping a database when drop database fails” in the “Other Useful Tasks” chapter in the *Troubleshooting and Disaster Recovery* guide.
- 2 Create a database for load. Make sure the database you create has sizes as least as large as those in `sysusages` for the original database (and that all other `sysusages` values match the original values). Refer to “create database” in the *Reference Manual: Commands* for more information on “create database for load.”

3 Load the database from backup (refer to “load database” in the *Reference Manual: Commands*).

4 Use the online database command to make the database available for use.

Version in which this error is raised

All versions

Error 625

Severity

21

Message text

Versions prior to 15.0:

Could not retrieve row from logical page %ld via RID because the entry in the offset table for that RID is less than or equal to 0.

Version 15.0 and later:

Adaptive Server failed to retrieve a row via its RID in database '%.*s' because the entry in the offset table for that RID is less than or equal to 0. %S_RID. %S_PAGE.

Note This error may be caused by a hardware problem.

Explanation

Every row in a database has a unique row ID (RID). The two parts of a row ID are a logical page number and a row number. The row number part of row IDs are stored at the end of every data page and indicate where a certain row is located on that page.

Commands that use row IDs include `dbcc checkdb`, `dbcc checktable`, and `dbcc checkcatalog`. They verify the consistency of indexes and use the row IDs stored in nonclustered indexes to quickly access the data rows.

Error 625 occurs when Adaptive Server tries to retrieve a row from a data page by specifying the row ID and it fails because the requested row ID has an illegal value (smaller than or equal to zero).

This error can occur under the following conditions:

- The data page of a nonclustered index is corrupt
- Hardware failure

Errors 624 and 625 occur under similar conditions.

Action Examine your operating system error log file and the Adaptive Server error log to determine if hardware errors may have corrupted your database devices.

Refer to “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for instructions about how to use the `dbcc page` command to determine the table and the index involved. Save this information and use it in the appropriate section that follows.

Warning! The procedure referenced in the preceding paragraph uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

The specific action you take depends on whether or not you know which index or query caused the problem. In general, dropping and re-creating the index will resolve the index corruption.

Index unknown, query known

If you do not know which index is causing the problem, but you do know which query encounters the problem, use this section to correct it. If you do not know the index or query, use the next section.

- 1 Determine which index should be dropped by reading the `showplan` output for the query that encounters the error and determining which nonclustered index the query is using to access the table in question.

Turn on `showplan` and use the `no execute` mode:

```
1> set showplan on
2> go
```

```
1> set noexec on
2> go
```

Enter the query that was causing the problem. For example:

```
1> select title from titles where title > "Cooking"
2> go
```

```
STEP 1
The type of query is SELECT
FROM TABLE
titles
Nested iteration
Index: titleind
```

Turn showplan and the no execute mode off again:

```
1> set noexec off
2> go

1> set showplan off
2> go
```

- 2 If the index identified in step 1 is on a system table (object ID is less than 100), refer to “Fixing a corrupted system table index” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for instructions on how to repair the system table index.
- 3 Otherwise, drop and re-create the index identified in step 1 (in this example “titleind”).

If this did not clear the error, create a clustered index on the table or drop and re-create the existing clustered index. Note that creating a clustered index requires an amount of space equal to 120 percent to 150 percent of the size of the table where you want to create the index.

Both index and query unknown

If you do not know either the index or the query, rebuild all nonclustered indexes on the table using the instructions in this section.

- 1 Look at the dbcc page output you obtained earlier.
- 2 If the object is a system table (object ID is less than 100), refer to “Fixing a corrupted system table index” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for instructions on repairing the system table index. Then go to step 4.
- 3 For a user table (object ID is 100 or greater), use sp_helpindex to list all indexes on the table, and then rebuild all the nonclustered indexes using one of two methods:

- Drop and re-create each nonclustered index on the table;

or:

- If a clustered index also exists on the table, drop and re-create it, which causes all nonclustered indexes to be automatically rebuilt. If your table is large, you may not have the space to do this (a rough rule of thumb is 120 percent to 150 percent of your table size must be available).

If the above methods did not clear the error, create a clustered index on the table (if space is available). You may then drop the clustered index.

- 4 Run `dbcc checktable` on the table to confirm that the problem has been resolved.

Additional information

For more information, refer to “create database” in the *Reference Manual: Commands* guide.

If you need to call Sybase Technical Support for assistance, have available the information listed in “Reporting errors” on page 4, including the output from `select * from master..sysusages`.

Version in which this error is raised

All versions

Error 629

Severity

21

Message text

```
Clustered index row entry for data page %ld is missing
in index page %ld of table '%S_OBJID' in database
'%S_DBID'; index row contains data page %ld instead. You
can recreate the clustered index to fix this error.
Xactid is (%ld,%d).
```

Explanation

This error occurs when Adaptive Server fails to delete a clustered index entry because the index entry did not point to the expected page.

In the error message text, the first page number refers to the data page and the last page number refers to the leaf-level index page that points to the data page.

Error 629 can occur when you attempt to delete a row in a table that has a clustered index, but Adaptive Server cannot find a pointer from the leaf page of the clustered index to the data page as expected.

The error is caused by data corruption that occurred during Adaptive Server processing (for example, an operating system panic occurs, causing interruption in disk writes when using UNIX files for Sybase database devices). This may be due to a problem with the server, the operating system, or hardware.

Action

- 1 Use the procedure in “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* to identify the table and index that correspond to the first page number in the error message text.

Warning! The procedure referenced in step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 If the object encountering the error is not a system table (a system table’s object ID is less than 100), continue with step 3.

If the object with the error is a system table, refer to “Fixing a corrupted system table index” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for instructions on how to repair the system table index. Then go to step 6.

- 3 Determine the index name and the first page in the page chain.

```
1> use <database_name>
2> go

1> select name, first from sysindexes
2> where id = <object_ID> and indid = 1
3> go
```

- 4 Check that the page linkage is intact.

```
1> dbcc pglinkage(<dbid>, <first page>, 0, 1, 0, 1)
2> go
```

If you see a message indicating “End of chain reached”, the page linkage is intact; continue with step 5. If you do not see such a message, the page linkage is bad; stop here, do not drop the clustered index, and contact Sybase Technical Support.

	<p>5 To ensure that the information needed to re-create the index is available, run <code>sp_helpindex</code> on the index prior to dropping it.</p> <p>Drop the index.</p> <p>Re-create the index. This clears the corruption in most cases.</p>
	<p>6 Run <code>dbcc checktable</code> on the table to verify that the corruption is gone. If corruption still exists, call Sybase Technical Support.</p>
Additional information	<p>If you need to call Sybase Technical Support for assistance, have available the information listed in “Reporting errors” on page 4.</p> <p>Refer to “drop index” and “create index” in the <i>Reference Manual: Commands</i> for information about dropping and creating indexes.</p>
Version in which this error is raised	All versions

Error 631

Severity	21
Message text	<p>Versions earlier than 15.0:</p> <p>The length of %d passed to delete row routine for the row at offset %d is incorrect on the following page: %S_PAGE. The expected row length is %d. The offset should be no greater than %d.</p> <p>Version 15.0 and later:</p> <p>The length of %d passed to delete row routine for the row at offset %d is incorrect on the following page of database '%S_DBID': %S_PAGE. The expected row length is %d. The offset should be no greater than %d.</p> <hr/> <p>Note This error may be caused by a hardware problem.</p> <hr/>
Explanation	<p>This error occurs when Adaptive Server attempts to delete a row (via a direct delete or inherently through updating) from an index or data page by specifying the row offset and the row length, and the action fails because the specified values of the offset or row length did not match the actual values.</p> <p>Error 631 can happen under the following conditions:</p>

- During normal processing, when Adaptive Server tries to delete the row specified by the error message.
- During database recovery. Database recovery occurs:
 - During Adaptive Server start-up
 - When a load database or load transaction command is processed

Some potential causes of error 631 are:

- Data corruption during normal processing (for example, an operating system panic occurs, causing interruption in disk writes when using UNIX files for Sybase database devices). This may be due to a problem with Adaptive Server, the operating system, or hardware.
- Hardware failure during loading or dumping.

Action

Error 631 is probably the result of a more serious underlying problem, and recovering from this error depends on when the error occurred. Follow the instructions in this section, selecting the correct set depending on whether the error occurred during normal processing or during database recovery.

If the error occurred during normal processing

- 1 Use the procedure in “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* to identify which table and index correspond to the page number from the error message text.

Warning! The procedure referenced in the step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 If the object encountering the error is not a system table (a system table’s object ID is less than 100), continue with step 3.

If the object with the error is a system table and the index ID is not 0, refer to “Fixing a corrupted system table index” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for instructions on how to repair the system table index.

If the index ID is 0, contact Sybase Technical Support. They may be able to help you repair the corruption, but you may have to restore from clean backups.

- 3 For user tables, if the index ID is 0 or 255, continue with step 4.

If the index ID is not 0 or 255, translate it into an index name:

```
1> use <database_name>
2> go

1> select name from sysindexes
2> where id = <object_ID> and indid = <index_ID>
3> go
```

To ensure that the information needed to re-create the index is available, run the `sp_helpindex` procedure on the index prior to dropping it.

Drop the index.

Re-create the index. This clears the corruption in most cases.

Run `dbcc checktable` on the table to verify that the corruption is gone.

- 4 If the index ID is 255, delete the bad data row.

If the index ID is 0, do one of the following:

- Restore the database from clean backups.
- Refer to “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide.

Warning! Some data might be lost on this page if you recover your table using `bcp` or `select into` (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

Before dumping your database, make sure it works correctly. Run the following commands prior to each dump:

```
1 dbcc checkdb.
```

- 2 dbcc checkalloc or dbcc checkalloc with the fix option. (Refer to “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for information about how to run these commands in multi-user mode and how to prevent spurious allocation errors from dbcc commands.)

Note Alternatively, using dbcc checkstorage with checkverify can examine for most types of consistency problems in less time.

If the error occurred during database recovery

When this error occurs during recovery, the database is marked suspect and is not accessible. Usually, you must load the database from backup. To do this, follow the instructions below:

- 1 Drop the database. If the drop fails, follow the instructions in “Dropping a database when drop database fails” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide.
- 2 Create a database for load. Make sure the database you create has sizes as least as large as those in sysusages for the original database (and that all other sysusages values match the original values). Refer to “create database” in *Reference Manual: Commands* for more information on “create database for load.”
- 3 Load the database from backup. Refer to “load database” in *Reference Manual: Commands*.)
- 4 Use the online database command to make the database available for use.

If loading from backups is not feasible, call Sybase Technical Support. Have the following documentation ready:

- Adaptive Server error log
- Text of all error messages
- `select * from master..sysusages output`

Preventing this error on recovery

One of the following items may help to prevent error 631 from occurring on recovery:

- Use shutdown to shut down Adaptive Server instead of shutdown with nowait after a period of heavy update activity in your databases;

or:

- checkpoint each database that is being used before shutting down Adaptive Server.

Examine both your operating system error log and the Adaptive Server error log to determine if hardware errors may have affected your database devices. Look for any kernel messages reporting I/O errors and check the hardware error log or diagnostics utilities for I/O errors.

Additional information

Refer to “Developing a Backup and Recovery Plan” in the *System Administration Guide: Volume 2* for complete information about how to safely create, dump, load, and re-create databases.

Version in which this error is raised

All versions

Error 644

Severity

21

Message text

```
Index row entry for data row id (%ld, %d) is missing
from index page %ld of index id %d of table '%S_OBJID'
in database '%S_DBID'. Xactid is (% ld,%d). Drop and
re-create the index.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when the nonclustered index indicated by “index id” is corrupt. This corruption is detected when a process tries to delete a nonexistent row.

Action

To recover from this error, use the following steps to drop and re-create the index.

- 1 Record the value of “index page” and “index ID” specified in the 644 error text. Then follow the instructions in “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* to identify which table and index correspond to the index page number. Also note the object ID.

- 2 If the object with the error is a system table (its object ID is less than 100), refer to “Fixing a corrupted system table index” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for instructions on how to repair the system table index. Then go to step 4.
- 3 If the object ID is greater than 100, drop and re-create the index, using the table name and the index name obtained in step 1. This clears the corruption in most cases.
- 4 To verify that all problems have been resolved on this table, run `dbcc checktable` and `dbcc tablealloc` on the affected table.

Additional information

If problems persist, this procedure may not be sufficient to clean up the index corruption, and you should contact Sybase Technical Support. Have the following information ready:

- Server version and SWR version level
- Server error log
- Output of `dbcc tablealloc` and `dbcc checktable`
- Text of all error messages

Version in which this error is raised

All versions

Error 678

Severity

20

Message text

Invalid Buffer Cache Access: Could not open object '%ld' in database '%d'

Explanation

This error is raised when Adaptive Server is unable to locate and read an object into a data cache. The error occurs when `dbcc checkalloc` finds that an allocation structure or extent belongs to a non-existent object.

The error does not prevent you from continuing other operations. However, each occurrence of this error represents a loss of up to eight pages on disk, or 16K. The pages cannot be used until the error is corrected.

Action

This error may be raised due to an Adaptive Server problem. Contact Sybase Technical support for assistance.

Version in which this error is raised All versions

Error 691

Severity 20

Message text Encountered invalid logical page '%ld' while accessing object '%ld' in database '%d'. This is an internal system error. Please contact Sybase Technical Support.

Explanation This error occurs when Adaptive Server tries to access an object but requests an invalid page number. There are two possibilities:

- A negative page number is requested.
- Page number 0 (zero) is requested. To be more specific, a request is issued for page 0, which is an allocation page, as if it was a data page. Data pages do not use page ID 0. Corrupt page links or incorrect entries for the object in `sysindexes` can cause the server to request page 0.

Error 691 results in a stack trace and breaks your connection to Adaptive Server.

Note These errors may be due to hardware problems.

Action Take the following steps:

- 1 Check the server error log for other errors that may have been raised prior to the 691 error. Save the entire error log.
- 2 Identify the object named in the error message using the following `isql` commands:


```
1> use <database name>
2> go

1> select name from sysobjects where id=<object id>
2> go
```
- 3 To determine the full extent of the corruption, run `dbcc checkstorage`, the `dbcc checkdb` and `dbcc checkalloc` commands, or `dbcc checktable` and `dbcc tablealloc` as soon as possible.

- 4 To check if this error is a result of hardware failure, examine your operating system error log and correct hardware problems.

If the 691 error is not a soft error (one that exists only in memory), the problem is severe. Sybase Technical Support may be able to help you recover from the error. However this recovery procedure often results in loss of data; if this is unacceptable, you will probably need to restore from backups.

Note If a Replication Agent™ is configured on this database, investigate replication solutions by contacting Sybase Technical Support.

Additional information	Refer to the write-up for “Error 605” on page 72 for a discussion of potential causes of hardware errors. Have the information available that is listed in “Reporting errors” on page 4 before calling Sybase Technical Support, including the output of the dbcc checks from step 3 above.
Version in which this error is raised	All versions

Error 692

Severity	20
Message text	Uninitialized logical page '%ld' was read while accessing object '%ld' in database '%d'. Please contact Sybase Technical Support.
Explanation	Error 692 occurs when Adaptive Server tries to read a page that is not formatted (initialized) or is improperly formatted. The server retries the read if the first read attempt failed, and raises the 692 error if the second read attempt also fails. Error 692 results in a stack trace and breaks your connection to Adaptive Server. <hr/> Note These errors may be due to hardware problems. <hr/>
Action	Take the following steps: <ol style="list-style-type: none">1 Check the server error log for other errors that may have been raised prior to the 692 error. Save the entire error log.

- 2 Obtain page information by running `dbcc page` as soon as possible, using the database ID and page number shown in the message:

```
1> dbcc page (<db_id>, <page_number>, 0, 1, 1, -1)
2> go
```

Warning! Use the `dbcc page` command only as directed above. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided here for reference, you should use this command only in the specific diagnostic situations described, and with the specific syntax shown in this section. Failure to do so could cause performance problems and/or database corruption.

- 3 Identify the object named in the error message using the following `isql` commands:

```
1> use <database name>
2> go
```

```
1> select name from sysobjects where id=<object_id>
2> go
```

- 4 Find information about this object in `sysindexes` (pre-15.0) or `syspartitions` (version 15.0 and later):

- Versions earlier than 15.0:

```
1> select first, root, doampg, ioampg from sysindexes
2> where name=object_name(<object_id>)
3> go
```

- Version 15.0 and later:

```
1> select firstpage, rootpage, datoampage, indoampage
2> from syspartitions
3> where id=<object_id>
4> go
```

- 5 To determine the full extent of the corruption, run `dbcc checkstorage`, the `dbcc checkdb` and `dbcc checkalloc` commands, or `dbcc checktable` and `dbcc tablealloc`) as soon as possible.

- 6 To check if this error is a result of hardware failure, examine your operating system error log and correct hardware problems. See “Checking the operating system error log” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for assistance.

Warning! If the 692 error is not a soft error (one that exists only in memory), the problem is severe. Sybase Technical Support may be able to help you recover from the error. However, this recovery procedure often results in loss of data; if this is unacceptable, you will probably need to restore from backups.

Additional information

Refer to the “Error 605” on page 72 write-up for a discussion of potential causes of hardware error. Refer to “dbcc” in the “Commands” chapter of the *Reference Manual: Commands* for more dbcc information.

Have the information on hand that is listed in “Reporting errors” on page 4 ready before calling Sybase Technical Support, including the output of the sysindexes or syspartitions query and dbcc checks from steps 2, 4, and 5 above.

A database dump and transaction log dumps may also be required.

Version in which this error is raised

All versions

Error 693

Severity

20

Message text

Multiple copies of logical page '%ld' from database '%d' reside in more than one cache. This is an internal system error. Please contact Sybase Technical Support.

Explanation

This error occurs when Adaptive Server tries to access an object but is unable to establish the identity of a logical page for the object. Copies of the page are found in more than one cache, and the current session is accessing the wrong version.

Error 693 is caused by an Adaptive Server problem. It results in a stack trace and breaks your connection to Adaptive Server.

Action

- 1 Check the server error log for other errors that may have been raised prior to the 693 error. Save the entire error log.

- 2 Obtain cache information by running `dbcc page` as soon as possible, using the database ID and page number shown in the message:

```
1> dbcc page (<db_id>, <page_number>, 0, 1, 1, -1)
2> go
```

Warning! Use the `dbcc page` command only as directed above. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided here for reference, you should use this command only in the specific diagnostic situations described, and with the specific syntax shown in this section. Failure to do so could cause performance problems and/or database corruption.

- 3 Contact Sybase Technical Support with the information you collected.

Additional information

Refer to “Error 605” on page 72 for a complete list of related object mismatch errors.

Have the information available that is listed in “Reporting errors” on page 4 before calling Sybase Technical Support, including the output of the `dbcc` checks.

Version in which this error is raised

All versions

Error 694

Severity

24

Message text

An attempt was made to read logical page '%ld', virtpage '%ld' from virtual device '%d' for object '%ld' in database '%d'. The page was not read successfully. You may have a device problem or an operating system problem.

Explanation

This error occurs when Adaptive Server tries to access an object but is unable to read the page mentioned in the message. Adaptive Server automatically issues a second read request to verify the consistency of the first read attempt. One of the following cases will be true:

- The first I/O request did not take place, and consequently the page header is unchanged by the I/O operation. Adaptive Server writes the following message into the error log:

```
I/O did not occur, buffer contents are unchanged.  
Previous page in buffer = %ld Previous objid = %ld
```

- The second read request was also unsuccessful or it returned incorrect information.

```
I/O did not occur, buffer contents are unchanged.  
Previous page  
in buffer = %ld Previous objid = %ld
```

- The second read request was successful. The first read request may have failed due to a timing or caching problem on the device.

In all cases, Adaptive Server writes the following message into the error log:

```
Suspect a Device or OS problem, %s
```

and raises error 694, resulting in a stack trace and breaking the connection to Adaptive Server.

This error indicates hardware or operating system problems.

Action

Take the following actions to resolve the problem:

- Save the entire Adaptive Server error log.
- Check whether dbcc checkalloc or dbcc checkdb is running concurrently with other activities on the server. Error 694 may occur when too many I/Os are requested from the server.
- Check if the device in question provides a caching mechanism. If so, turn off caching to see if this resolves the 694 error.
- If the error persists, examine your operating system error log and correct any hardware problems. See “Checking the operating system error log” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for assistance.

Note On the IBM Regatta platform, 694 errors may be encountered when you perform I/O operations on the IBM Regatta family of servers, which is caused by a cache line synchronization problem on this family of servers. See TechNote 1023360 at <http://www.sybase.com/detail?id=1023360> for details and the solution from IBM.

Additional information	Refer to the write-up for “Error 605” on page 72 for a discussion of potential causes of hardware errors.
Version in which this error is raised	All versions

Error 695

Severity	20
Message text	<p>Versions prior to 15.0:</p> <p>An attempt was made to read logical page '%ld' for object '%ld' in database '%d' from cache '%.*s'. Wrong logical page '%ld' was brought into cache.</p> <p>Version 15.0 and later:</p> <p>An attempt was made to read logical page '%ld' for %S_PTNINFO from disk. Wrong logical page '%ld' was brought into cache '%.*s'.</p>

Note These errors may be due to hardware problems.

Explanation	<p>Error 695 occurs when Adaptive Server tries to read a page from disk (the first logical page in the message), but the page number in the header of the returned page does not match the requested page. This means that the page is corrupt on disk, and has been overwritten by another page. The error may be due to an Adaptive Server problem but could also be caused by problems such as overlapping partitions or hardware errors.</p> <p>Error 695 results in a stack trace and breaks your connection to Adaptive Server. Additional errors may be reported on the console and in the server error log prior to this error.</p>
Action	<ol style="list-style-type: none"> 1 Check the server error log for other errors that may have been raised prior to the 695 error. Save the entire error log. 2 Obtain page information by running <code>dbcc page</code> as soon as possible, using the database ID and page number shown in the message: <pre>1> dbcc page (<db_id>, <page_number>, 0, 1, 1, -1)</pre>

2> go

Warning! Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 3 To determine the scope of the problem, locate the device on which the database resides, and check if any other databases use the same device. Refer to “Determining the physical device on which a database resides” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.
- 4 Check that the device is partitioned correctly. See “Choosing between raw partitions and UNIX files” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.
- 5 Run `dbcc log` to find the history of transactions against this page.

```
1> dbcc log (<dbid>, 0, <page_no>, 0, 0, -1)
2> go
```

Warning! Use the `dbcc log` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 6 To determine the full extent of the corruption, run `dbcc checkstorage`, the `dbcc checkdb` and `dbcc checkalloc` commands, or `dbcc checktable` and `dbcc tablealloc` as soon as possible.
- 7 To check if this error is a result of hardware failure, examine your operating system error log and correct hardware problems. See “Checking the operating system error log” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for assistance. Also see “Error 605” on page 72 for a discussion of potential causes of hardware errors.

Additional information	<p>Sybase Technical Support may be able to help you recover from the error. However this recovery procedure often results in loss of data; if this is unacceptable, you will probably need to restore from backups.</p> <p>Refer to “dbcc” in the “Commands” chapter of the <i>Reference Manual: Commands</i> for more dbcc information.</p> <p>Have the information listed in “Reporting errors” on page 4 before calling Sybase Technical Support, including the output of the dbcc checks from steps 2, 5, and 6 above.</p>
Version in which this error is raised	All versions

Error 696

Severity	21
Message text	<p>Versions prior to 15.0:</p> <pre>An attempt was made to fetch logical page '%ld' in tempdb from cache '%.*s'. Page belongs to object '%ld' and not to object '%ld'. Restart of SQL Server will clear the error. Please contact your System Administrator for help.</pre> <p>Version 15.0 and later:</p> <pre>An attempt was made to fetch logical page '%ld' from cache '%.*s'. Page belongs to %S_PTNINFO and not to %S_PTNINFO. Restart of ASE will clear the error. Please contact your System Administrator for help.</pre>
Explanation	<p>This error occurs when Adaptive Server discovers page allocation corruption. Adaptive Server tries to access a particular object but discovers a page in the object's page chain whose object ID is different than that of the object being accessed.</p> <p>This error is similar to the 605 error. However, error 696 always occurs in the context of the tempdb database. Since the object ID mismatch is detected in cache, this is a transient (soft) error.</p>
Action	Shut down and restart Adaptive Server to clear the cache problem. If error 696 occurs again, call Sybase Technical Support.

Additional information	Have the following information ready before calling Sybase Technical Support: <ul style="list-style-type: none">• Adaptive Server version and SWR version level• Text of all error messages• Text of the query which raises the error.
Version in which this error is raised	All versions

Error 697

Severity	20
Message text	<pre>An attempt was made to fetch logical page '%ld' for object '%ld' in database '%d' from cache '%.*s'. Wrong logical page '%ld' was found in cache.</pre>
<hr/> Note These errors can be due to hardware problems. <hr/>	
Explanation	<p>This error occurs when Adaptive Server tries to access an object in cache but arrives at an invalid page. The page number is correct on disk, but it does not match the page number of the page in memory.</p> <p>Error 697 results in a stack trace and breaks your connection to Adaptive Server.</p>
Action	Since the problem occurs in cache, it is a soft error. Shut down and restart Adaptive Server to clear the memory corruption.
Version in which this error is raised	All versions

Memory Manager Errors (700s)

This section contains error messages for the Adaptive Server Memory Manager.

Error	Page
Error 701	107
Error 702	110
Error 703	111
Error 706	113
Error 707	113
Error 709	114
Error 714	115
Error 715	116
Error 717	116
Error 718	117
Error 719	117
Error 720	118
Error 721	119

Error 701

Severity

17

Message text

There is not enough procedure cache to run this procedure, trigger, or SQL batch. Retry later, or ask your SA to reconfigure SQL Server with more procedure cache.

Explanation

This error occurs when the procedure cache is not large enough to execute a stored procedure, trigger, or batch of one or more SQL statements. This can happen when loading a query plan into procedure cache or when resolving or compiling a procedure, trigger, or batch.

The amount of available procedure cache is based on Adaptive Server use, so this error may occur intermittently as the procedure cache content is used and released. Also, one or more procedures can repeatedly fail with this error while others complete normally. This is because query plans vary greatly in size. Therefore, while there may not be sufficient space to load a large query plan, several smaller plans may fit.

Index creation uses sort buffers to hold data during sort operations. Each sort buffer reserves procedure cache, in proportion to the number of rows on the data page in the buffer; if the data rows are short, there are many rows per page, requiring more cache space to be allocated per page. This factor, multiplied by the number of sort buffers needed, can also exhaust procedure cache leading to the 701 error. For corrective action in this scenario, see “701 error raised during index creation” on page 110.

Error 701 is raised with the following states:

State	Meaning
1	Unable to read a procedure into cache from sysprocedures.
2	Unable to allocate memory for a procedure header.
3	Unable to find sufficient space to accommodate the query in the current procedure header.
4	Unable to allocate memory for a procedure header based on existing page allocation.
5	Unable to allocate space from kernel memory.
7	Unable to write a procedure from cache into sysprocedures.
8	Similar to state 1.
9	Unable to allocate memory during a log scan.
10	Unable to allocate memory for a sort.
11	Unable to allocate memory for a procedure header.
12	Unable to allocate memory for statement cache.
13	Unable to allocate memory during a dump transaction command using the "with standby access" option.
14	Unable to allocate memory for parsing client SQL.
15	Unable to allocate memory for a bitstring operation.
16	Unable to allocate memory for a bitstring operation.
17, 18, 19	Used only by the Adaptive Server Enterprise exerciser.
20	Unable to allocate memory during an alter table.

Action

Version 12.5 and later

Correct this error by increasing the configuration parameter procedure cache size to make more procedure cache available.

If the error is seen during index creation, see the title “701 error during index creation” in this section for corrective action.

Version pre-12.5

Correct this error by increasing the size of procedure cache. You can increase the amount of total procedure cache in three ways:

- Use the `total memory` configuration parameter to increase the total amount of memory requested by Adaptive Server.
- Use the `procedure cache percent` configuration parameter to increase the percentage of total cache space dedicated to procedure cache.
- Decrease the amount of memory required for other resources; for example, unneeded user connections.

Increasing the `total memory` configuration parameter for Adaptive Server is the most straightforward method. It increases both the procedure and data caches, although it can waste memory.

Note On most UNIX systems, be sure to make any necessary operating system memory resource adjustments. For example, verify the current kernel value for the maximum size of a shared memory segment, usually `SHMMAX`, and make sure you have adequate memory and swap space on the system for the additional memory configured for the Server. For more information, refer to a) your operating system documentation, and b) the Adaptive Server installation and configuration guide.

If your memory resources are limited, you can increase the `procedure cache percent` configuration parameter for the Adaptive Server without changing the value of the `total memory` configuration parameter. This shifts space from the data cache to the procedure cache and may result in performance degradation or other problems if not enough data cache remains.

You can also increase the size of the procedure cache without substantially changing the amount of data cache. This requires some combination of increasing both the `total memory` and `procedure cache percent` configuration parameters in such a way that the majority of the new memory goes to procedure cache. Although this method requires more planning, it allows you to control where the additional memory goes. See the sections entitled “Memory Use and Performance” in *Performance and Tuning: Basics* and “Configuring Memory” in *System Administration Guide: Volume 2* for more information about configuring Adaptive Server memory.

701 error raised during index creation

If 701 error is raised when creating an index:

- Increase procedure cache, or
- Decrease the number of sort buffers using `sp_configure`, or
- Decrease the number of rows per page using `max_rows_per_page`, or by increasing the row size with dummy columns.

Version in which this error is raised

All versions

Error 702

Severity

20

Message text

Memory request for %d bytes exceeds the size of single page of %d bytes.

Explanation

This error occurs when you exceed any of the following limits:

- 128 search conditions or join operations in a SQL statement.

A search condition sets the conditions in a `where` or `having` clause. For more information about, and examples of, search conditions, see the sections “Queries: Selecting Data from a Table” and “Using Aggregates, Grouping, and Sorting” in the *Transact-SQL User’s Guide*.

A join operation compares two or more tables or views. For more information about and examples of join operations, see the section entitled “Joins: Retrieving Data from Several Tables” in the *Transact-SQL User’s Guide*.
- The total width of the columns in a temporary table during a union exceeds the allowed row width (on a server of 2K page size, rows consist of 1962 characters; allow 2 characters of row overhead for all-pages-locked (APL) tables, for a maximum usable row size of 1960. For data-only-locked (DOL) tables, subtract two characters per varchar column in determining usable row size).

After displaying error 702, Adaptive Server terminates the current process.

Action

If your query exceeds the limit of 128 search conditions or join operations, rewrite the query so that the limit of 128 statements is not exceeded.

If the total width of the columns in a temporary table during a union exceeds the allowed row width, rewrite your query so that the limit (1962 bytes on a 2K page server) is not exceeded.

For example, if a table has a unique key on column *<column1>*, run the following query in order to delete the row in the table that contains the unique key *<unique_key1>*:

```
1> begin transaction
2> delete <table_name>
3> where <column1> = <unique_key1>
4> go
```

If the key is unique in the table, only one row will be deleted by the above query. If only one row is being returned by the above query, commit the transaction with the following query:

```
1> commit transaction
2> go
```

Otherwise, roll back the transaction:

```
1> rollback transaction
2> go
```

If you are not sure if the key *<unique_key1>* is unique in *<column1>*, you can check by running the following query:

```
1> select * from <table_name>
2> where <column1> = <unique_key1>
3> go
```

Similarly, if the table has a unique index on columns *<column1>* and *<column2>*, you can delete a row in that table by running the following query:

```
1> delete <table_name>
2> where <column1> = <unique_key1>
3> and <column2> = <unique_key2>
4> go
```

Version in which this error is raised

All versions

Error 703

Severity

17

Message text	You cannot run this procedure, trigger, or SQL batch because it requires more than %ld pages of memory. Break it up into shorter queries, if possible.
Explanation	<p>This error occurs when a stored procedure or trigger cannot be executed because it requires more memory than is allowed for execution.</p> <p>Two different stages of execution can trigger this error: “resolution,” in which the query tree is built, or “compilation,” in which the query plan is generated. If a query has been executed successfully but later fails with this error, it means that the query tree used to be less than the allowed memory limit but has since grown. If a query encounters this error every time it is executed, then it is too complex to be executed in its present form without exceeding the memory limit. Refer to <i>Performance and Tuning Series: Query Processing and Abstract Plans</i> for detailed information about query trees and query plans.</p>
Action	<p>Errors 701 and 703 are very similar. Error 703 is raised when the procedure requires multiple process headers (control structures) but there is not enough memory to allocate the next header; error 701 is raised when there is not enough memory to extend the current header which already has some pages allocated. See “Error 701” on page 107 for additional troubleshooting information.</p> <p>If the procedure or trigger causes this error every time you try to execute it, divide it into smaller pieces. This division process varies greatly with the type of SQL statement.</p> <p>If the object has successfully executed before, the 703 error probably occurred because the query tree grew beyond the specified page limit. This growth occurs each time the query tree is re-resolved. Once you have encountered the 703 error in this situation, you can drop and re-create the procedure or trigger in order to shrink the query tree to a legal size, or you can break the object up into smaller modules. Restarting Adaptive Server has no effect on the size of the query tree because query trees are stored on disk.</p> <p>If this error occurs frequently on an object, you can either periodically drop and re-create the object as part of regular database maintenance, or break it into smaller modules.</p>
Version in which this error is raised	All versions

Error 706

Severity	20
Message text	Process %d tried to remove PROC_HDR 0x%lx that it does not hold in Pss.
Explanation	<p>The Memory Manager allocates, deallocates, and manages memory for Adaptive Server. It manages an array of structures in the procedure header, each of which represents a physical page of memory in procedure cache and the current byte allocation of that page.</p> <p>As part of its memory deallocation process, Adaptive Server stops tracking procedure headers when they are no longer needed and attempts to remove them from procedure cache. Error 706 occurs when Adaptive Server fails to deallocate a procedure header.</p> <p>Error 706 is caused by memory corruption or an Adaptive Server problem.</p>
Action	Since procedure headers are stored in procedure cache, restarting Adaptive Server should clear error 706. If it does not clear the 706 error or if the 706 error occurs again, contact Sybase Technical Support.
Version in which this error is raised	All versions

Error 707

Severity	20
Message text	System error detected during attempt to free memory at address 0x%lx. Please consult the SQL Server error log for more details.
Explanation	<p>The Memory Manager allocates, deallocates, and manages memory for Adaptive Server. It manages an array of structures in the procedure header, each of which represents a physical page of memory in procedure cache and the current byte allocation of that page.</p> <p>As part of its memory deallocation process, Adaptive Server tries to release the pages of memory allocated to a procedure header when they are no longer needed. When Adaptive Server is unable to release that section of memory, error 707 occurs.</p> <p>Error 707 is caused by memory corruption or an Adaptive Server problem.</p>

Action	Since procedure headers are stored in procedure cache, restarting Adaptive Server should clear error 707. If it does not clear the 707 error or if the 707 error occurs again, contact Sybase Technical Support.
Version in which this error is raised	All versions

Error 709

Severity	17
Message text	<code>There is insufficient system memory to continue login process for spid %d.</code>
Explanation	<p>The Memory Manager allocates and deallocates memory for Adaptive Server processes and manages memory requirements for the system. For each process connecting to the server, the Memory Manager allocates a memory structure called a procedure header and ensures that there is enough memory to dedicate to the process.</p> <p>Error 709 is raised when a user process attempts to log in to the server, but there is not enough memory available to complete the login.</p>
Action	<p>Correct this error by increasing the amount of available memory. You can do this by:</p> <p><i>Version 12.0 and earlier:</i></p> <ul style="list-style-type: none">• Use the <code>total memory</code> configuration parameter to increase the total amount of memory requested by Adaptive Server.• Use the <code>procedure cache percent</code> configuration parameter to increase the percentage of total cache space dedicated to procedure cache.• Decrease the amount of memory required for other resources; for example, unneeded user connections. <p>Increasing the <code>total memory</code> configuration parameter for Adaptive Server is the most straightforward method. It increases both the procedure and data caches, although it can waste memory. Since <code>total memory</code> and <code>procedure cache percent</code> are static parameters, you must restart Adaptive Server after changing these parameters.</p>

Version 12.5 and later:

- Use the max memory configuration parameter to increase the total amount of memory available to Adaptive Server.
- Use the procedure cache size configuration parameter to increase the percentage of memory dedicated to procedure cache.
- Decrease the amount of memory required for other resources; for example, unneeded user connections.

Increasing the procedure cache size configuration parameter for Adaptive Server is the most straightforward method.

Additional information

Check the value of the number of user connections configuration parameter. Since there is memory overhead associated with each user connection, setting this parameter too high can reduce the amount of space available for the data and procedure caches, and contribute to 709 errors.

Version in which this error is raised

All versions

Error 714

Severity

20

Message text

Unable to locate a procedure cache memory control descriptor for memory `0x%p` allocated by process %d.

Explanation

The procedure cache memory control descriptor keeps track of memory allocated from the procedure cache. The memory allocated could be a procedure header (proc header) or just a chunk of memory. During process cleanup, either due to successful process completion or due to some earlier error, Adaptive Server will raise this message when it is not able to find the procedure cache memory control descriptor for a procedure cache structure or for the memory chunk specific to this task, that was being freed.

Action

There is no action to be performed in the event of this error. In a worse case scenario, there will be an incorrect procedure cache memory accounting.

Additional information

A configurable shared memory dump (CSMD) configured for error 714 may be helpful in determining root cause for this error. Information about CSMD can be found in the section “Other Useful Tasks” in the *Troubleshooting and Disaster Recovery* guide.

Version in which this error is raised Version 15.0 and later

Error 715

Severity	20
Message text	Unable to locate a procedure cache memory control descriptor for memory '0x%p' in the global list.
Explanation	For each memory structure allocated in procedure cache, there is a procedure cache memory control descriptor that is used to track procedure cache usage. During process cleanup, either due to successful process completion or due to some earlier error, Adaptive Server Enterprise was unable to find the procedure cache memory control descriptor for a global memory structure in procedure cache that was being freed.
Action	There is no action to be performed in the event of this error. In a worse case scenerio, there will be an incorrect procedure cache memory accounting.
Additional information	A configurable shared memory dump (CSMD) configured for error 715 may be helpful in determining root cause for this error. Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 717

Severity	20
Message text	The last prochdr associated with the passed in main prochdr is incorrect. Last prochdr: `0x%p`. Main prochdr: `0x%p`. This is an Adaptive Server internal error. Please report the problem to Sybase Technical Support.
Explanation	When allocating memory for a process, Adaptive Server Enterprise starts the search for available memory with the last prochdr. While checking this address against the main prochdr it detects an inconsistency, and reports the 717 error

Action	This is an internal inconsistency. Please contact Sybase Technical Support
Additional information	A configurable shared memory dump (CSMD) configured for error 717 may be helpful in determining root cause for this error. Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 718

Severity	20
Message text	<p>Memory allocation from process private memory is being performed by an incorrect process. Current Spid: '%d'. Owning Spid: '%d'. Memory: '0x%p'.</p> <p>This is an Adaptive Server internal error. Please contact Sybase Tech Support.</p>
Explanation	Before allocating a new PROC_HDR, Adaptive Server searches existing procdhrs to attempt to find free space from which to satisfy a memory request. In this search it encounters a PROC_HDR that is neither a global PROC_HDR (rare) nor owned by the spid requesting memory. Adaptive Server Enterprise should not be examining PROC_HDRs owned by another spid, so it raises this 718 error.
Action	This is an internal Adaptive Server Enterprise error and there is no action that can be taken. Contact Sybase Technical Support.
Additional information	A configurable shared memory dump (CSMD) on error 718 may be required to reach root cause. Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 719

Severity	20
----------	----

Message text	The last prochr value is not available for the prochr '0x%p'. This is an Adaptive Server internal error. Please contact Sybase Tech Support.
Explanation	<p>Adaptive Server Enterprise keeps the address of the last PROC_HDR in the process memory control structure and uses it as a memory allocation hint because it is likely to have free space. If this address is invalid, Adaptive Server Enterprise raises error 719.</p> <p>Error 719 is associated with four states:</p> <ul style="list-style-type: none">• When Adaptive Server Enterprise is allocating memory, the address of the last PROC_HDR is NULL in the process memory control structure.• When merging PROC_HDR chains, one main PROC_HDR becomes secondary to another main PROC_HDR, the target. The address of last PROC_HDR of the secondary PROC_HDR is NULL.• When merging PROC_HDR chains, one main PROC_HDR becomes secondary to another main PROC_HDR, the target. The address of last PROC_HDR of the target PROC_HDR is NULL.• When moving a subordinate PROC_HDR from one principal PROC_HDR to another principal PROC_HDR, the target, the address of last PROC_HDR of the target PROC_HDR is NULL.
Action	This is an internal Adaptive Server Enterprise error and there is no action that can be taken. Contact Sybase Technical Support.
Additional information	A configurable shared memory dump (CSMD) on error 719 may be required to reach root cause. Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 720

Severity	17
Message text	Unable to allocate a procedure memory control descriptor. Please increase the value of the configuration parameter '%s' or try again when there is less activity on the system.

Explanation	For each memory structure allocated in procedure cache, whether it is for a PROC_HDR or just a chunk of procedure cache, there is a procedure cache memory control descriptor that is used to track procedure cache usage. Adaptive Server Enterprise has attempted to allocate the space for a procedure cache memory control descriptor from the procedure cache memory control pool and the allocation failed.
Action	Usually the configuration parameter listed will be procedure cache size. Increasing the procedure cache size may allow a larger procedure cache memory control pool.
Additional information	If the MDA tables are in use, monitoring memory allocations may reveal all allocations prior to the one that failed. If after increasing the procedure cache size, this error continues to occur, a configurable shared memory dump (CSMD) on error 720 may assist in determining root cause. Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 721

Severity	20
Message text	The prochr '0x%p' has large allocation enabled incorrectly. This is an Adaptive Server internal error. Please contact Sybase Tech Support.
Explanation	<p>“Large allocations” is an internal Adaptive Server Enterprise algorithm designed to improve the speed of memory allocations by attempting to allocate more than one page of memory at a time.</p> <p>Adaptive Server Enterprise will not store procedures to disk with the large allocation algorithm enabled, because loading the same procedures back into memory may not always be possible. Therefore, the procedure headers used for stored objects must not use large allocation.</p> <p>Error 721 is raised when Adaptive Server attempts to write a procedure to disk and large allocations is enabled.</p>
Action	If you encounter error 721 when trying to run a create procedure command, retry that command.

Error 721

Additional information	If error 721 persists, a configurable shared memory dump (CSMD) may assist in determining root cause. Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0.1 and later

Buffer Manager Errors (800s)

The Adaptive Server Enterprise Buffer Manager works in conjunction with the Cache Manager to manage physical disk pages while they are in memory. This section contains error messages for the Buffer Manager.

Error	Page
Error 803	121
Error 804	122
Error 806	123
Error 813	125
Error 820	127
Error 821	127
Error 822	129
Error 823	131
Error 832	132
Error 834	133
Error 835	135
Error 840	135
Error 842	136
Error 847	137
Error 849	138
Error 852	139
Error 855	141
Error 861	142
Error 863	143
Error 881	143
Error 890	144

Error 803

Severity

20

Message text

Unable to place buffer '0x%lx' from cache '%.*s'

holding logical page '%ld' in sdes for object '%.*s' - either there is no room in sdes or buffer already in requested slot.

Note “sdes” is an abbreviation for “session descriptors.” Refer to “Error 603” on page 71 for information about session descriptors.

Explanation

A session descriptor is an internal data structure in Adaptive Server which contains information about a table. All open objects in Adaptive Server require a descriptor structure. Whenever a connection to Adaptive Server attempts to open a table, it gets a session descriptor.

To ensure that a buffer read by the session remains in the buffer cache and does not age out, Adaptive Server “keeps” the buffer. The session descriptor contains eight slots where buffers are stored that are “kept” by the session.

Error 803 occurs during recovery or regular Adaptive Server operation when one of the following is true:

- Adaptive Server runs out of session descriptors
- A buffer was already in the session descriptor’s requested slot

Action

When an 803 error occurs:

- Look for other errors in the Adaptive Server error log to indicate the specific source of the problem (for example, not enough locks, not enough open objects).
- Increase the appropriate configuration parameter.

If problems persist, contact Sybase Technical Support for assistance. Before contacting Sybase Technical Support, be prepared to provide the complete text of all error messages. Examine your operating system error log file as well as the Adaptive Server error log to determine if hardware errors may have caused the problem.

Version in which this error is raised

All versions

Error 804

Severity

20

Message text

Unable to find buffer 0x%lx holding logical page %ld in

```
sdes 0x%lx kept buffer pool for object '%.*s'.
```

Note “sdes” is an abbreviation for “session descriptors.” Refer to “Error 603” on page 71 for information about session descriptors.

Explanation	<p>A session descriptor is an internal data structure in Adaptive Server that contains information about a table. All open objects in Adaptive Server require a descriptor structure. Whenever a connection to Adaptive Server attempts to open a table, it gets a session descriptor.</p> <p>To ensure that a buffer read by the session remains in the buffer cache and does not age out, Adaptive Server “keeps” the buffer. The session descriptor contains eight slots where buffers that are “kept” by the session are stored.</p> <p>Error 804 occurs when Adaptive Server cannot find the pointer to a buffer header in a session descriptor. This error can be transient.</p>
Action	When an 804 error occurs, restart Adaptive Server as soon as possible. If additional 804 errors occur, call Sybase Technical Support.
Additional information	Before contacting Sybase Technical Support, be prepared to provide the complete text of all error messages and have the information on hand listed in “Reporting errors” on page 4.
Version in which this error is raised	All versions

Error 806

Severity	21
Message text	Could not find virtual page for logical page %ld in database '%S_DBID'.

Note This error may be caused by a hardware problem.

Explanation	A virtual page is a page within a Sybase device. A logical page is a page in an Adaptive Server database. There is a one-to-one correspondence between these two types of pages.
-------------	--

Error 806 occurs when Adaptive Server fails to convert a logical page number to a virtual page number. Depending on what caused the error, it can be serious or transient.

If error 806 is accompanied by the message “Page %d is not in the range of pages for database id %d,” it means you used `dbcc page` with an invalid parameter. This is not a serious problem.

If error 806 occurs on recovery, it may be transient or serious (see information under “Action” for specifics).

Error 806 can occur during normal processing, such as creating an index or running a stored procedure. In this case, the error is probably caused by corruption or a problem with Adaptive Server and it is a serious error.

Action

If error 806 specifies `tempdb` in the message output, restart Adaptive Server. Since `tempdb` is rebuilt each time Adaptive Server is restarted, this may clear the error. If the error occurs again (on `tempdb`), call Sybase Technical Support.

During recovery

If error 806 occurs on recovery, the database will be marked suspect. If the error is transient, resetting the suspect status will solve the problem. To resolve this problem:

- 1 Bypass recovery by starting Adaptive Server with status `-32768`.
- 2 Run `dbcc checkdb` and `dbcc checkalloc` on the database listed in the error message output.
- 3 Set status back to 0.
- 4 Shut down Adaptive Server:

```
1> shutdown with nowait
2> go
```
- 5 Try recovery again.

If error 806 occurs again, call Sybase Technical Support to help you recover from this error if the corruption is not too widespread. However, recovery from clean backups may be necessary.

During normal processing

Run `dbcc checkdb` and `dbcc checkalloc` on the database listed in the error message output. Call Sybase Technical Support. Technical Support may be able to help you recover from this error if the corruption is not too widespread. However, recovery from clean backups may be necessary.

Additional information	<p>Refer to the chapter “Checking Database Consistency” in the <i>System Administration Guide: Volume 2</i> for information about dbcc commands.</p> <p>Before contacting Sybase Technical Support, have the information on hand listed in “Reporting errors” on page 4, including dbcc checkdb, dbcc checkalloc, and dbcc page output.</p>
Version in which this error is raised	All versions

Error 813

Severity	20
Message text	Logical page '%ld' in database '%S_DBID', cache '%.*s' is already hashed.

Note This error may be caused by a hardware problem.

Explanation	<p>This error may be serious, especially if it occurs on a table’s data page, which means that a page is currently in the data cache and is in use by a table or index but is not marked as allocated. An attempt to allocate it raises error 813.</p> <p>The same page could be allocated again after it is removed from the data cache, resulting in a loss of whatever data resides on the page.</p>
-------------	---

Warning! Pages encountering error 813 will not be included in a database dump. This is because database dumps are performed by reading allocation pages and not by traversing page chains. Therefore, this error should be corrected before dumping the database.

After the page is removed from the data cache, further attempts to access this page may raise 2500 series errors (allocation errors) when running dbcc checkalloc, dbcc tablealloc, or dbcc indexalloc.

Action	<p>To clear the error:</p> <ol style="list-style-type: none"> 1 Execute the checkpoint command on the database specified in the error message. This will cause modified pages in the data cache to be flushed to disk.
--------	---

- 2 Run `dbcc checkalloc` or `dbcc checkalloc` with the `fix` option on the database. As a result, you might get allocation errors because the page displayed in the error message will be linked but not allocated. Refer to “`dbcc`” in *Reference Manual: Commands* for information about running these commands.
- 3 If you do get allocation errors:
 - Check your hardware error log or diagnostics utilities for I/O errors.
 - Check your operating system error log file and the Adaptive Server error log to determine if hardware errors may have caused the corruption. (Look for messages reporting I/O errors.)
 - Refer to the write-ups in this manual.
- 4 If you still have 813 errors after fixing any other errors reported during step 2, shut down and restart Adaptive Server. This will probably clear error 813.

If errors still occur, call Sybase Technical Support. They might be able to help you recover from this error if the corruption is not too widespread. However, recovery from backups might be necessary.

Using dump transaction with `no_log`

Using the `dump transaction with no_log` command can result in an 813 error. Therefore, do not use `dump transaction with no_log` unless it is absolutely necessary (when `dump transaction with truncate_only` will not truncate the log). Try using `dump transaction with truncate_only` first, then use the `no_log` option only as a last resort. Refer to “Error 1105” on page 173 for details.

Check for any occurrences of `dump transaction with no_log` in any of your scripts or procedures and replace them with `dump transaction with truncate_only`.

Additional information

For more information on this series of errors, refer to the information in this section and in “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.

Version in which this error is raised

All versions

Error 820

Severity	21
Message text	Attempt to dirty non-log buffer %S_BUF which is in I/O.
Explanation	

Adaptive Server uses buffers to manage physical pages while they are in memory. When a page is read into memory from the disk, a buffer header is assigned to it. The header tracks the usage and contents of the page image.

The buffers themselves are grouped together to form Memory Address Space Segments (MASSes). A buffer cache consists of MASSes linked in MRU/LRU (most recently used/least recently used) chains.

A buffer that was changed while in cache and has not yet been written to disk is known as a dirty buffer.

Error 820 is raised when the server detects that an invalid operation has been performed on a buffer or MASS. The error is raised in the following states:

State	Meaning
1	Attempted to mark as dirty a MASS that is currently being written.
2	Object ID mismatch detected when linking a buffer to the chain.
3	MASS found to belong to a different database. This state is only raised by diagserver.
4	Buffer is being changed but the page residing in the buffer does not belong to the transaction log as expected.
5	Buffer pinning to a cache is invalid with respect to transaction state.
6	Attempted to pin a buffer to a cache while the buffer is pinned to a different cache.

Some scenarios in which this error may be raised include:

- Parallel select into with union all.
- drop role when 15 or more databases exist on the server.

Action	Call Sybase Technical Support.
Version in which this error is raised	All versions

Error 821

Severity	20
----------	----

Message text Attempt to unhash buffer at 0x%lx with a buffer pageno of %ld and database id %ld with HASHED status set failed - buffer was not found. %S_PAGE.

Note This error may be caused by a hardware problem.

Explanation Every data or index page in the data cache has a corresponding buffer pointing to it, controlling its stay in the cache. As these pages are read and modified, they are scheduled to be written to disk and removed from the data cache. Error 821 occurs when Adaptive Server attempts to remove a nonexistent or corrupted page from the data cache.

Some potential causes of this error are:

- Overlapping partitions (refer to “Choosing between raw partitions and UNIX files” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*).
- Hardware failure.
- An Adaptive Server problem (such as can occur when using the dump transaction with `no_log` command after a period of heavy user activity in the database).

This error can occur during database recovery as well as during normal operation of Adaptive Server.

Action 1 Check the Adaptive Server error log to determine whether there are other indications of hardware problems, such as kernel messages reporting I/O errors. If problems exist, solve them.

 2 Check the operating system error log or run diagnostic utilities to check for I/O errors. If problems exist, solve them.

 Refer to “Checking the operating system error log” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for assistance.

 3 Once the actual cause of corruption has been removed, shut down and restart Adaptive Server to clear the 821 error.

Using dump transaction with no_log

Using the dump transaction with `no_log` command can result in an 821 error. Therefore, do not use dump transaction with `no_log` unless it is absolutely necessary (when dump transaction with `truncate_only` will not truncate the log). Try using dump transaction with `truncate_only` first and use the `no_log` option only as a last resort. Refer to Error 1105 for details.

Check for any occurrences of dump transaction with `no_log` in any of your scripts or procedures and replace them with `dump transaction with truncate_only`.

Version in which this error is raised

All versions

Error 822

Severity

20

Message text

Could not start I/O for request BLKIO flags = 0x%lx, size = %ld, errcode = 0x%lx, %S_BUF.

Note This error may be caused by a hardware problem, usually a missing or offline device.

Explanation

This error occurs when a read or write operation was requested and Adaptive Server encountered an error while attempting to initiate the request. Error 822 can occur for any of the following reasons:

- A database device is offline
- A database device or file has been removed or renamed
- A database device or file is inaccessible to Adaptive Server for any other reason, such as insufficient permissions or a device number higher than the number of devices configured for Adaptive Server (such a device will not be activated at startup)

The information in the error message refers to an internal Sybase structure and does not help determine which database device or file is involved. Other error messages appearing in the Adaptive Server error log just before the 822 error usually indicate the nature of the problem and the device involved. Refer to “Additional Information” for examples of this type of error message.

Action

- 1 Examine the availability and condition of the device involved using your standard operating system procedures, and make sure the device is accessible.

- 2 Adaptive Server will not be able to recover the database upon restarting because an 822 error will mark the database suspect. If you know that the database was marked suspect because the device was unavailable, reset the status of the database to allow recovery to continue.

Warning! Do not use these procedures if other errors in the error log near the 822 error message suggest that the database may have been marked suspect for some other reason.

Reset the suspect status using one of the methods supplied in “Resetting a database’s “suspect” status” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery*. After you reset the status, and start Adaptive Server, if the database is again marked suspect, examine the Adaptive Server error log for indications of other errors.

Additional information Here are two examples of the 822 error and associated messages from the Adaptive Server error log. Examining the error messages just before the error usually helps determine the type of problem.

Example 1: permission problems

In this example, the permissions on *disk1* (virtual device 6) were incorrect, so Adaptive Server was unable to access this device:

```
kernel: initializing virtual device 6, "/work/disk1"  
kernel: dopen: open "/work/disk1", Permission denied  
kernel: dinit: failed to open primary device /work/disk1 for vdn6  
kernel: udstartio: vdn 6 has not been set up  
server: Error: 822, Severity: 20, State: 3  
server: Could not start I/O for request BLKIO ....
```

To correct an error like this:

- 1 Change the ownership to user “sybase”. Make sure that this user has read and write permissions on the device. In our example, check the user's permission on */work/disk1*.
- 2 Become the “sybase” user.
- 3 Start Adaptive Server.

Example 2: file missing

In this example, the UNIX file corresponding to *disk1* (virtual device 6) did not exist when Adaptive Server started, causing access to the device to fail:

```
kernel: initializing virtual device 6, "/work/disk1"  
kernel: dopen: open "/work/disk1", No such file or directory  
kernel: dinit: failed to open primary device /work/disk1 for vdn6
```

```
kernel: udstartio: vdn 6 has not been set up
server: Error: 822, Severity: 20, State: 3
server: Could not start I/O for request BLKIO ....
```

To correct an error like this, make sure that the virtual device in the error message exists, and correct whatever error caused Adaptive Server not to find it. (For example, it may have been renamed or moved while Adaptive Server was shut down.)

If the virtual device no longer exists, you will have to restore from known clean backups. Merely creating an empty device with the right name will not solve this problem.

Version in which this error is raised

All versions

Error 823

Severity

24

Message text

I/O error detected during %S_MSG for %S_BUF.

Note This error may be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server encounters an I/O error on a read or write request made to a Sybase device. It usually means you have disk problems. The parameters in the error message refer to internal Sybase structures and do not often help determine which device is involved. However, additional kernel messages in Adaptive Server's error log recorded before the 823 error was raised should indicate which device is involved.

Causes of the 823 error can include the following:

- If the sum of `vstart` and `size` exceeds the device size. This can happen on AIX systems but is also possible on other UNIX platforms if `sysusages` or `sysdevices` have been manually altered.
- When using UNIX files as devices, if the file system is full but the device appears to have space as tracked by `sysusages` and `sysdevices`. This raises the 823 error during `create database` or `alter database`.

Action Check the accessibility and condition of the device in question. Once you have identified and corrected the problem, execute the `dbcc checkdb` and `dbcc checkalloc` commands to ensure that no other damage was caused by the bad device.

AIX systems only

Error 823 may occur if the disk device is run under the logical volume manager (LVM), `vstart` is equal to 2 and the size is not shortened by 2. `vstart` needs to be equal to 2 because the first AIX block (1 AIX block equals 2 Sybase pages) has to be available for the LVM control block (LVCB).

Note You may encounter false 823 errors when you perform I/O operations on the IBM Regatta family of servers. This is caused by a cache line synchronization problem on this family of servers. See TechNote 1023360 at <http://www.sybase.com/detail?id=1023360> for details and the solution from IBM.

Also, on AIX 5.3 false 823 errors can occur on any IBM servers (not just Regattas) due to IBM APAR IY80355, which may show up on any version that does not include the fix in Adaptive Server Enterprise version 12.5.2 ESD #2, version 12.5.3 ESD #1, and version 15.0.x).

Version in which this error is raised All versions

Error 832

Severity 20

Message text Unable to place buffer 0x%x holding logical page %ld in sdes for object '%.*s' - there is no room in sdes aux keep pool.

Explanation Error 832 is raised when the server runs out of space in the 2K sort buffer during a sort-intensive operation such as a parallel index creation.

Error 832 is raised with the following states:

State	Meaning
1	Buffer could not be placed in auxiliary buffer keep pool.

State	Meaning
2	Similar to state 1, but this state occurs when an operation needs to “keep” an already kept buffer.

Action

State 1

Increase number of sort buffers beyond the default value. This allows the sort manager to utilize more sort buffers and reduce the number of merge runs and amount of disk I/O required for the operation.

State 2

If this error is raised with state 2, report the problem to Sybase Technical Support.

Additional information

The number of sort buffers configuration parameter specifies the number of 2K buffers used to hold pages read from input tables and perform index merges during sorts.

Sybase recommends that you leave this parameter set to the default except when you are creating indexes in parallel. Setting the value too high can rob non-sorting processes of access to the 2K buffer pool in caches being used to perform sorts.

For more information on configuring number of sort buffers for parallel create index statements, see “Configuration Parameters” in the chapter “Setting Configuration Parameters” in the *System Administration Guide: Volume 1*. Also refer to “Configuring enough sort buffers” in the chapter “Indexing for Performance” in the Adaptive Server Enterprise version 15.0 or later *Performance and Tuning: Locking* guide.

Version in which this error is raised

All versions.

Error 834

Severity

20

Message text

Illegal attempt to clean buffer: %S_BUF.

Explanation The Cache Manager works in conjunction with the Buffer Manager to manage physical disk pages while they are in memory. When a page is read into memory from the disk, a buffer header is assigned to it. The header tracks the usage and the contents of the page image. The page and the buffer header together define a buffer.

The buffers themselves are grouped together to form memory address space segments (MASSes). A buffer cache consists of MASSes linked in MRU/LRU (most recently used or least recently used) chains. The Cache Manager manages MASSes and buffer caches.

When a buffer or MASS is no longer needed (for example during a drop database, drop table, drop index, or deallocation of pages), Adaptive Server performs some cleanup tasks and releases the resources taken up by the buffer or MASS. Before these actions are carried out, Adaptive Server does a number of checks to make sure the buffer or MASS is no longer in use. Error 834 occurs when one of these checks fails. It is caused by an Adaptive Server problem.

Error 834 occurs with the following states:

State	Meaning
1	If the MASS is actively being used or if Adaptive Server is in the process of writing out the MASS when Adaptive Server tries to clean it up, error 834 occurs with state 1. This state is only raised by diagserver.
2	If the buffer is still hashed or has not been cleaned up when Adaptive Server is ready to release it, error 834 occurs with state 2. This state is only raised by diagserver.
3	If the buffer is actively being used or if Adaptive Server is in the process of writing out the buffer when Adaptive Server is ready to release it, error 834 occurs with state 3. This state is only raised by diagserver.
4	If the buffer belonging to a database has been set aside for later use and Adaptive Server tries to clean it out of a cache, error 834 occurs with state 4.
5	If a buffer belonging to an object has been set aside for later use and Adaptive Server tries to clean it out of a cache, error 834 occurs with state 5.

Action

State 4, LTMs running

834 errors with state 4 can occur if the log transfer manager (LTM), a Replication Server component, is running during server recovery. The LTM process attempts to log in before the database is fully recovered. To avoid this error, shut down the LTM before restarting the server. Restart LTM only after all the databases have been recovered.

All other states

Shut down and restart Adaptive Server to clear the buffer structure. If error 834 occurs again, call Sybase Technical Support.

Version in which this error is raised All versions

Error 835

Severity 20

Message text `Illegal attempt to change contents of buffer: %S_BUF.`

Explanation
 The Cache Manager works in conjunction with the Buffer Manager to manage physical disk pages while they are in memory. When a page is read into memory from the disk, a buffer header is assigned to it. The header tracks the usage and the contents of the page image. The page and the buffer header together define a buffer.

When a client process completes, Adaptive Server performs some cleanup tasks such as closing the buffers and releasing the resources taken up by the buffers. If the client process terminates abnormally, however (for example if the process is killed during execution), Adaptive Server may be unable to carry out the appropriate cleanup, buffers are left open, and error 835 is raised. The error can also occur due to an Adaptive Server problem, when a page being updated is deallocated before the modification can complete.

Action
 If the problem re-occurs, shut down and restart Adaptive Server to clear the buffer structure. If error 835 continues to occur, call Sybase Technical Support.

Version in which this error is raised All versions

Error 840

Severity 17

Message text
 Device '%.*s' (with physical name '%.*s', and virtual device number %d) has not been correctly activated at startup time. Please contact a user with System Administrator (SA) role.

Note This error may be caused by a hardware problem.

Explanation	<p>This error can occur when Adaptive Server is unable to access a device during startup.</p> <p>This error can occur for any of the following reasons:</p> <ul style="list-style-type: none">• A database device is offline• A database device or file has been removed or renamed• A database device or file is inaccessible to Adaptive Server for any other reason, such as insufficient permissions
Action	<ol style="list-style-type: none">1 Examine the availability and condition of the device involved using your standard operating system procedures, and make sure the device is accessible.2 Adaptive Server cannot recover databases that have space on the device upon restarting because error 840 marks the databases as suspect. If you know that the databases were marked as suspect because the device was unavailable, resolve the problem with the device and reset the status of the databases to allow recovery to continue. <hr/> <p>Warning! Do not use these procedures if other errors in the error log near the 840 error message suggest that the databases may have been marked suspect for some other reason.</p> <hr/>
Version in which this error is raised	<p>Reset the suspect status using one of the methods supplied in “Resetting a database’s “suspect” status” in the chapter “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i>.</p> <p>After you reset the status, execute the <code>dbcc checkdb</code> and <code>dbcc checkalloc</code> commands for the affected databases to ensure that no other damage was caused by the device being unavailable. If other errors occur, refer to the write-ups for those errors in this manual. If errors still persist, call Sybase Technical Support.</p> <p>All versions</p>

Error 842

Severity	20
Message text	Attempt to set bufclaims in PSS structure for process

	<code>%d</code> to negative value while unclaiming cache.
Explanation	<p>Adaptive Server manages physical disk pages while they are in memory by maintaining buffer pools of logical pages in cache. It allocates buffers to server tasks as needed, keeps track of the buffers assigned to caches and tasks, and releases buffers that are no longer needed by a task. It also accesses the process status structure (PSS), which is a memory structure associated with each process describing the state of the process.</p> <p>Error 842 occurs when a buffer is released, and the task is found to have a negative buffer count as a result. This error is at times seen during sort operations; for example, during index creation.</p> <p>Error 842 breaks your connection to Adaptive Server.</p>
Action	<p>Reconnect to the server.</p> <p>If the error persists:</p> <ol style="list-style-type: none"> 1 Configure the server to obtain a shared memory dump for future occurrences of error 842; and 2 Contact Sybase Technical Support after collecting the memory dump and related error data. <p>To configure the server to dump shared memory upon future occurrence of error 842:</p> <pre>1> sp_configure "dump on conditions",1 2> go 1> sp_shmdumpconfig "add", error, 842, 1, <directory>, <file_name> 2> go</pre>
Additional information	Have the information listed in “Reporting errors” on page 4 on hand when you call Sybase Technical Support, including the output of <code>sp_shmdumpconfig</code> .
Version in which this error is raised	All versions

Error 847

Severity	10
Message text	<pre>Checkpoint process detected hardware error writing logical page '%ld', virtual page '%ld' for dbid '%ld', cache '%.*s'. It will sleep until write completes</pre>

successfully.

Note This error may be caused by a hardware problem.

Explanation

Adaptive Server's automatic checkpoint mechanism guarantees that data pages changed by completed transactions are regularly written from memory to the database device. To accomplish this, the checkpoint process periodically checks the number of log records in the transaction log for each database. The number of log records, and a recovery interval defined with `sp_configure`, are used to decide when to checkpoint each database. During the checkpoint all pages that have been modified in memory, but not on disk, since the last checkpoint are written out to the database device. `checkpoint` is performed on a per database basis for all the databases on an Adaptive Server.

Error 847 occurs when the checkpoint process is unable to write out pages to the database device due to a hardware problem. The error is preceded by the following message in the error log:

```
bufwritedes: write error detected - spid=%ld, ppage=%ld,
bvirtpg=%ld, dbid=%ld
```

Action

This is an informational message and the checkpoint task will periodically retry writing pages to the device. To identify the affected devices so you can investigate the hardware problem, refer to the chapter "Creating and Managing User Databases" in the *System Administration Guide: Volume 2* and read the section "Getting Information About Database Storage."

Note If hardware problems were detected, execute the `dbcc checkdb` and `dbcc checkalloc` commands for the affected database to ensure that no database corruption occurred due to the problem.

Version in which this error is raised

All versions

Error 849

Severity

24

Message text

I/O error detected for device '%.*s' (with physical

name '%.*s', and virtual device number %d).

Note This error may be caused by a hardware problem.

Explanation	<p>This error occurs when Adaptive Server encounters an I/O error on a read or write request made to a Sybase device. The error is accompanied by error 823 and is seen during dbcc operations.</p> <p>The parameters in the error message refer to the Sybase logical device name, the physical operating system location of the device, and the virtual device number (a unique identifier for the device).</p> <p>Error 849 is related to:</p> <ul style="list-style-type: none"> • Hardware problems on a disk where a Sybase device resides. • A disk mirroring problem in Adaptive Server. Disk mirrors allow an Adaptive Server database device to be duplicated, that is, all writes to the device are also copied to a separate physical device. If an I/O to a mirrored device fails, Adaptive Server “unmirrors” the bad device and continues to run unmirrored. Due to the disk mirroring problem, while a disk failure in the primary device unmirrors the device, it does not stop the server from initiating further I/O’s to the unmirrored device, raising error 849.
Action	<p>Check your operating system logs for any information about hardware problems on a disk containing a Sybase device. Because error 849 accompanies and provides supporting data for 823 errors, refer to “Error 823” on page 131 for more troubleshooting information. In the event of hardware failure, you will need to restore the affected databases from clean backups.</p> <p>If your site uses disk mirroring, error 849 may indicate that Adaptive Server continued to issue I/O’s to a failed primary device. Contact Sybase Technical Support for assistance.</p>
Additional information	<p>When calling Sybase Technical Support, have available the information listed in “Reporting errors” on page 4.</p>
Version in which this error is raised	<p>All versions</p>

Error 852

Severity 20

Message text Command Failed: Invalid request to move '%ld' buffers from the '%ld'K pool of the '%.*s' cache to the '%ld'K pool. Source pool only contains '%ld' buffers. Retry the command specifying fewer buffers.

Explanation The Cache Manager works in conjunction with the Buffer Manager to manage physical disk pages while they are in memory. When a page is read into memory from the disk, a buffer header is assigned to it. The header tracks the usage and the contents of the page image. The page and the buffer header together define a buffer.

The buffers themselves are grouped together to form memory address space segments (MASSes). A buffer cache consists of MASSes linked in MRU/LRU (most recently used/least recently used) chains. The Cache Manager manages MASSes and buffer caches.

The Adaptive Server command `sp_poolconfig` allows you to create, drop, resize, and get information about memory pools within named data caches. Error 852 occurs when you use `sp_poolconfig` or a configuration file to create or resize memory pool size and Adaptive Server is unable to find contiguous memory to fulfill the request.

Error 852 occurs with the following states:

State	Meaning
1	Error 852 occurs with state 1 when the size of the destination pool is smaller than the requested memory size.
2	Error 852 occurs with state 2 when the requested pool size is larger than the source pool size.

Action 1 Use `sp_cacheconfig` to display information about the current configuration of the affected data cache and memory pools in the cache. For example:

```
1> sp_cacheconfig test_cache
2> go
```

Cache Name	Status	Type	IO Sz	Wash Sz	Config Value	Run Value
test_cache	Active	Mixed			2.00 Mb	2.00 Mb
test_cache	Active	Mixed	2 Kb	512 Kb	0.00 Mb	1.50 Mb
test_cache	Active	Mixed	4 Kb	100 Kb	0.50 Mb	0.50 Mb
Total					2.00 Mb	2.00 Mb

- The first line identifies the cache. Lines after the first line show defined memory pools.
- The IO Sz column shows the I/O size for a memory pool.

	<ul style="list-style-type: none"> • The Config Value column shows the size the cache or pool will have after the next time Adaptive Server is restarted. If the value is 0, the size has not been explicitly configured and a default value will be used. • The Run Value column shows the size of the cache or pool now in use on Adaptive Server.
	2 Try the <code>sp_poolconfig</code> command again with appropriately-sized values or use <code>sp_poolconfig</code> to reconfigure the source pools to have an appropriate number of buffers available.
Additional information	Refer to the <i>Reference Manual: Procedures</i> for information about <code>sp_cacheconfig</code> and <code>sp_poolconfig</code> .
Version in which this error is raised	All versions

Error 855

Severity	16
Message text	Bad prefetch size of %d encountered.
Explanation	<p>On select, update, and delete statements, you can specify a prefetch size that specifies the I/O size in kilobytes for tables bound to caches for which large I/Os are configured. Valid values for size are 2K, 4K, 8K, and 16K. You must have configured the I/O size for the cache before you can specify it for a prefetch size. The procedure <code>sp_helpcache</code> shows the valid sizes for the cache to which an object is bound, or for the default cache.</p> <p>When you specify a prefetch size, the Adaptive Server parser passes that value to the Buffer Manager. If the value you specified is not available, error 7380 (a warning) is raised, and processing continues with the prefetch size that the Adaptive Server optimizer determines to be the “best” value.</p> <p>Error 855 occurs when:</p> <ul style="list-style-type: none"> • A problem occurs with the passing of the prefetch value and • The prefetch size that is passed to the Buffer Manager is smaller than the minimum allowed size or larger than the maximum allowed size.
Action	Call Sybase Technical Support.

Version in which this error is raised All versions

Error 861

Severity 16

Message text Command Failed: Cannot bind object '%ld', dbid '%d' to cache because you have an open cursor on the target object. Retry the command after closing the cursor.

Explanation You can use the stored procedure `sp_bindcache` to bind a database, table, index, or text or image object to a named data cache. If there is an open cursor on the target object, the cursor structures have an open descriptor for the object and there may be cache information stored there. Error 861 occurs when you use `sp_bindcache` or the configuration file to change object binding and Adaptive Server finds that the object has been opened by a cursor.

Action Close the cursor and try the command again:

- 1 Get information about cursors you have open:

```
1> sp_cursorinfo null
2> go
```

- 2 If the output from step 1 shows an open cursor for the object you are trying to bind, close the cursor:

```
1> close <cursor_name>
2> go
```

- 3 If the output from step 1 did not show an open cursor for the object you are trying to bind, this probably means someone else opened it. Use `sp_lock` to determine who has locks on the object. If you can determine who has the open cursor, ask them to close it.

- 4 If you cannot determine who has the open cursor on the object you are trying to bind, shut down and restart Adaptive Server or wait and try again later.

- 5 Try your command again.

Version in which this error is raised All versions

Error 863

Severity 20

Message text Buffer resources in cache %s, id %d are unavailable.
Please re-run this query or ask the system administrator to re-configure buffer cache memory.

Explanation When a page is read into memory from disk, a buffer header is assigned to it. The header tracks the usage and the contents of the page image. The page and the buffer header together define a buffer. The buffers themselves are grouped together to form memory address space segments (MASSes). A buffer cache consists of MASSes linked in MRU/LRU (most recently used/least recently used) chains. Error 863 occurs when you run a query and Adaptive Server is unable to find an empty buffer in any buffer pool in the named cache.

Error 863 occurs with the following states:

State	Meaning
1	Adaptive Server could not find a buffer pool that had any available buffers in it.
2	When Adaptive Server makes a prefetch request for a MASS of a specific size, it checks the page range for the request to determine whether any of the pages within the range are already in cache. If any of the pages are already in cache, error 863 occurs with state 2. This state is caused by an Adaptive Server internal error.

Action Rerun the query later (when buffer resources in that cache might be free) or use `sp_cacheconfig` to increase the size of the cache.

Additional information Refer to the most recent version of the *Reference Manual: Procedures* for information about `sp_cacheconfig`.

Version in which this error is raised All versions

Error 881

Severity 20

Message text Could not start I/O for request for a page in extended cache BLKIO flags = 0x%1!, size = %2!, errcode = 0x%3!, %4!.

Explanation	Adaptive Server cannot issue an I/O from a buffer that is in an extended cache. This error is similar to 822, which is raised when Adaptive Server cannot start an I/O on a buffer in a standard cache. See the write-up for “Error 822” on page 129 for more details on possible causes and actions.
Action	Contact Sybase technical support to report this problem. If this error persists, A configurable shared memory dump (CSMD) configured may be helpful in determining root cause for this error.
Additional information	Information about CSMD can be found in the chapter “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.5.2 and later

Error 890

Severity	20
Message text	<pre>The mass containing the buffer '%S_BUF' should not be written in its current state. This is an Adaptive Server internal error. Please contact Sybase Technical Support.</pre>
Explanation	Rarely, Adaptive Server Enterprise may detect a breakage of synchronization between a process releasing a disk buffer in cache and one writing the buffer. This message indicates that Adaptive Server did not write this buffer to avoid a potential corruption.
Action	Contact technical support, to report this problem. If the error persists, a configurable shared memory dump (CSMD) may be helpful in determining the root cause for this error.
Additional information	Information about CSMD can be found in the chapter “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide
Version in which this error is raised	Version 12.5 and later

Open Database Manager Errors (900s)

This section contains error messages for the Open Database Manager.

Most 9xx errors are encountered during recovery. Many of these errors are raised as a result of the master database being corrupt. If this is the case, refer to the most recent version of the *Troubleshooting and Disaster Recovery* guide for assistance.

Error	Page
Error 903	145
Error 905	147
Error 906	149
Error 908	149
Error 911	150
Error 913	152
Error 916	153
Error 921	155
Error 924	156
Error 925	157
Error 926	160
Error 930	161
Error 935	163
Error 940	164
Error 941	166
Error 945	167
Error 950	168
Error 965	169

Error 903

Severity

23

Message text Unable to find row in sysindexes for clustered index on system catalog %d in database %d. This index should exist in all databases. Run DBCC CHECKTABLE on sysindexes in the database.

Note This error can be caused by a hardware problem.

Explanation This error occurs when Adaptive Server is unable to access the sysindexes table for a database.

Error 903 can occur under the following circumstances when Adaptive Server accesses sysindexes:

- Actions that involve opening a database since those actions access the sysindexes table for that database (for example, database recovery)
- Actions that need to access columns in sysindexes. Such actions include:
 - order_by clause compilation.
 - Running the update statistics command.
 - Threshold Manager space calculations. Refer to the chapter “Managing Free Space with Thresholds” in the *System Administration Guide: Volume 2*.

Possible causes of error 903 are:

- A problem has occurred with the device on which sysindexes resides and sysindexes has become corrupted.
- The partition on which the database resides has been mounted as a file system or overlapping partitions exist. Refer to “Choosing between raw partitions and UNIX files” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide.
- An Adaptive Server problem.

Action Determine whether the device on which the database resides is damaged and solve any problems that exist.

If the database named in the error message is accessible, run dbcc checkdb and dbcc checkcatalog to determine whether the database is corrupted. If the database is corrupted, restore it from backup.

If you cannot find any device problems and the database is not corrupted, call Sybase Technical Support.

Additional information	<p>Before calling Technical Support, have the information listed in “Reporting errors” on page 4 available, including the output from:</p> <ul style="list-style-type: none"> • <code>select * from <database_name>..sysindexes where id=object_id(“sysindexes”), and</code> • <code>dbcc checkdb and dbcc checkcatalog.</code>
Version in which this error is raised	All versions

Error 905

Severity	17
Message text	Unable to allocate a DBTABLE descriptor to open database ‘%S_DBID’. Close or drop another database before opening this one, or ask your System Administrator to raise the configuration parameter ‘number of open databases’.
Explanation	<p>As part of open database processing, the Database Table Manager controls and synchronizes access to database tables. When Adaptive Server tries to open a database and there are not enough database descriptors available, error 905 occurs. (The total number of database descriptors available is controlled by the number of open databases configuration parameter.)</p> <p>The action in the text of the 905 message only applies to environments where you cannot have any more open databases (for instance, because of memory limitations).</p>
Action	<p>When error 905 occurs, select one of the following strategies to solve the problem.</p> <p>Short term</p> <p>In the short term, you can close or drop another database. This is not a long-term solution. You can use this option until an Adaptive Server restart can be scheduled, as a restart is necessary for the long-term solution.</p> <p>To close a database, make sure no users are accessing the database or put the database into single-user mode. This does not guarantee that the database will not be accessed by Adaptive Server (for example, for an automatic checkpoint), so the 905 error may still occur.</p>

Long term

As a long-term solution, use `sp_configure` to increase the value of the number of open databases configuration parameter:

- 1 Determine your current value:

```
1> sp_configure "number of open databases"  
2> go
```

- 2 Determine what your new value should be. One way to do this is to determine the number of databases you have now and then add a padding factor for future growth.

For example:

```
1> select count(*) from sysdatabases  
2> go  
-----  
12
```

- 3 Set the new value. For example:

```
1> sp_configure "number of open databases", 15  
2> go
```

where “15” is the new value (12 existing databases plus a padding factor of 3 for future growth).

- 4 This step is not required for version 15.0 and later.

If you have a version of Adaptive Server Enterprise earlier than version 15.0, shut down and restart Adaptive Server to have the new value put into effect.

Note If you increase the value of the number of open databases configuration parameter, reconfigure Adaptive Server memory appropriately. 17,408 bytes are required for each additional open database. Refer to the chapter “Configuring Memory” in the *System Administration Guide: Volume 2* and the Adaptive Server Enterprise performance and tuning guides for information about Adaptive Server memory use.

Version in which this error is raised

All versions

Error 906

Severity	23
Message text	Could not locate row in sysobjects for system catalog %d in database %d. This system catalog should exist in all databases. Run DBCC CHECKTABLE on sysindexes in the database.

Note This error can be caused by a hardware problem.

Explanation Adaptive Server expects to find certain information in static locations within a database or on the server:

- Page 1 of each database is the first page for `sysobjects`
- Page 24 of each database is the first page for `sysindexes`

As part of open database processing, the Descriptor Manager searches page 1 for entries for `sysobjects` and page 24 for entries for `sysindexes`. If the entries are not found on the expected page, error 906 occurs.

Error 906 occurs during open database actions such as recovery and create database. Other possible causes are:

- The partition on which the database resides has been mounted as a file system or overlapping partitions exist. Refer to “Choosing between raw partitions and UNIX files” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.
- A database device has been created that exceeds the operating system maximum size limit (for example, 2GB on some 32-bit machines).

Action Error 906 means that, for the database specified in the message, either page 1 or page 24 is corrupted. Restore that database from backup.

Version in which this error is raised All versions

Error 908

Severity	22
Message text	Unable to find any entries in Sysusages for dbid '%d', database '%.*s'. Run DBCC CHECKTABLE on Sysusages

in the master database.

Explanation

The master..sysusages table keeps track of the space that you assign to Adaptive Server databases. create database and alter database allocate new space to the database by adding a row to sysusages for each database device or device fragment.

The Database Manager maps logical pages (pages in an Adaptive Server database) to virtual pages (pages in a database device). There is a one-to-one correspondence between these two types of pages.

Error 908 occurs during recovery when Adaptive Server attempts to generate the mapping of logical to virtual pages for a specified database, but is unable to find any rows for the database in the sysusages table.

Action

Check the sysusages table:

```
1> use master
2> go
```

```
1> dbcc checktable(sysusages)
2> go
```

If errors are reported, try to fix them using the sections in this manual that describe those errors.

If dbcc checktable reported no errors, and you do not need to recover the database in question, remove the database from your Adaptive Server installation. For further information, refer to “Dropping a database when drop database fails” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*. After removing the database, restart Adaptive Server and check that the 908 error is eliminated.

If dbcc checktable reported no errors, but you need to recover the database in question, start by removing the database from your Adaptive Server installation as described above. Next, rebuild the database using your creation scripts, and reload it from a clean backup.

Version in which this error is raised

All versions

Error 911

Severity

16

Message text	<p>Attempt to locate entry in sysdatabases for database '%.*s' by name failed - no entry found under that name. Make sure that name is entered properly.</p> <hr/> <p>Note This error can be caused by a hardware problem.</p> <hr/>
Explanation	<p>This error occurs during a use command when Adaptive Server tries to access a database that is not listed in sysdatabases. For example:</p> <ul style="list-style-type: none"> • A user tries to access a database not listed in sysdatabases: <pre>1> use <database_name> 2> go</pre> • A user tries to connect to Adaptive Server when an invalid database is indicated by the defaultdb column of master..syslogins. The user is connected only to master, or not connected at all, and error 911 occurs.
Action	<ol style="list-style-type: none"> 1 Obtain a list of all databases in sysdatabases by executing sp_helpdb or the following: <pre>1> select name from master..sysdatabases 2> go</pre> <p>The list returned contains the only valid databases recognized by Adaptive Server and available for use.</p> 2 Either create the missing database or change the reference to a database from the list created in step 1: <ul style="list-style-type: none"> • If a user-issued use command created the error, either create a database following regular procedures (refer to the chapter “Creating and Managing User Databases” in <i>System Administration Guide: Volume 2</i>) or change the use command to refer to one of the databases from the list created in step 1. • If the defaultdb option of sp_addlogin refers to an invalid database, the System Administrator (“sa”) or user needs to specify a new, valid, default database. To specify a new default database, choose a database name from the list in step 1 and use the following procedure: <pre>1> sp_modifylogin <user_name>, defdb, <database_name> 2> go</pre> <p>Although you can successfully change your own default database at any time with sp_modifylogin, you can access the database only if the database owner (“dbo”) has added you as a user in the database.</p>

Additional information	Refer to the <i>Reference Manual: Procedures</i> for information about <code>sp_addlogin</code> and <code>sp_modifylogin</code> .
Version in which this error is raised	All versions

Error 913

Severity	22
Message text	Could not find row in sysdatabases with database id %d. Run DBCC CHECKTABLE on sysdatabases.
Explanation	<p>This error occurs when Adaptive Server cannot find an entry in <code>sysdatabases</code> when looking for a specific database ID.</p> <p>Possible causes of error 913 are:</p> <ul style="list-style-type: none">• Accessing a stored procedure or view that refers to a table in a database that has been dropped.• Accessing a stored procedure or view that refers to a table in a database that has been dropped and re-created.• A nonclustered index on the <code>sysdatabases</code> table in the master database has been corrupted.• Running the system stored procedure <code>sp_help_rep_agent</code> without parameters, which is due to a problem in Adaptive Server Enterprise versions 12.0.x and version 12.5.0.2 or earlier. <p>This is a serious error if the index is corrupt or recovery fails.</p>
Action	<p>If the 913 error is caused by accessing a stored procedure or view that refers to a table in a database that has been dropped, either re-create the database or update the procedure or view to point to the appropriate database.</p> <p>When a database has been dropped and re-created, it can end up with a different database ID. If a stored procedure still references the old database ID, drop and re-create the stored procedure so that it will reference the new database ID. Refer to <i>drop procedure</i> and <i>create procedure</i> in the <i>Reference Manual: Commands</i> for information about dropping and re-creating stored procedures.</p> <p>If the cause of the error is a corrupted nonclustered index on a system table, use one of these options:</p> <ul style="list-style-type: none">• Use the <code>sp_fixindex</code> stored procedure to repair the index.

- Restore master from backup. See the chapter “Adaptive Server System Recovery” in most recent version of the *Troubleshooting and Disaster Recovery* guide for information about restoring the master database.

If the error was raised when you ran the system stored procedure `sp_help_rep_agent` without parameters, contact Sybase Technical Support for instructions on how to resolve this error.

Version in which this error is raised

All versions

Error 916

Severity	14
Message text	Server user id %d is not a valid user in database '%.*s'
Explanation	<p>When a user tries to open a database, Adaptive Server performs the following checks to determine the validity of the user:</p> <ol style="list-style-type: none"> 1 Looks for the <code>suid</code> (Adaptive Server user ID, copied from <code>syslogins</code>) of the process in <code>sysusers</code> to find the corresponding <code>uid</code> (user ID, unique in this database) and <code>gid</code> (group ID to which this user belongs). 2 Looks for a “guest” row in <code>sysusers</code> (the <code>suid</code> for “guest” is -1). 3 Looks for the user ID in <code>sysalternates</code>. 4 Checks <code>sysusers</code> to determine whether the user has “sso” role if the user does have an <code>suid</code> in the database. This check is related to Adaptive Server auditing. <p>Error 916 occurs when a user tries to open a database and Adaptive Server finds that the user ID is not valid.</p>
Action	<p>Decide whether you want the user to access the database with:</p> <ul style="list-style-type: none"> • A normal login and user ID • A guest user ID • An alias to another user ID • If you want the user to access the database with a normal login and user ID, determine whether the Adaptive Server user ID exists in the current database:

```
1> use <database_name>
2> go

1> select * from sysusers
2> where suid = <Server_user_ID>
3> go
```

Where *<Server_user_ID>* is the Adaptive Server user ID displayed in the error message. If the user is not found, use `sp_adduser` to add them.

If the user is found and you are getting error 916 for that user, run `dbcc checkdb` on the database to find out whether it is corrupt. If `dbcc checkdb` returns error messages, corruption has occurred. Follow the directions in this manual for recovering from those errors.

- If you want the user to access the database as `guest`, determine whether there is a “guest” row in `sysusers` of the database being opened:

```
1> use <database_name>
2> go

1> select * from sysusers
2> where suid = -1
3> go
```

Where *<database_name>* is the name of the database in the error message. If `guest` is not found, use `sp_adduser` to add it.

If `guest` is found, run `dbcc checkdb` on the database the user is trying to access to find out whether it is corrupt. If `dbcc checkdb` returns error messages, corruption has occurred. Follow the directions in this manual for recovering from those errors.

- If you want the user to access the database with an alias to another user, determine whether the user ID exists in `sysalternates` of the database being opened:

```
1> use <database_name>
2> go

1> select * from sysalternates
2> where suid = <Server_user_ID>
3> go
```

If the user ID is not found, use `sp_addalias` to add an alias for them.

If the user ID is found and you are getting error 916 for that user, run `dbcc checkdb` on the database the user is trying to access to find out whether it is corrupt. If `dbcc checkdb` returns error messages, corruption has occurred. Follow the directions in this manual for recovering from those errors.

- If the database is being accessed due to a referential constraint, make sure the user ID exists in each database.

If error 916 occurs after the above actions have been taken, call Sybase Technical Support.

Additional information

Refer to *Reference Manual: Procedures* for information about `sp_addlogin`, `sp_adduser`, and `sp_addalias`.

Have the information listed in “Reporting errors” on page 4 ready when you call Sybase Technical Support, including the output from:

- `syslogins`
- `sysusers` and `sysalternates` for the database being opened

Version in which this error is raised

All versions

Error 921

Severity

14

Message text

Database '%.*s' has not been recovered yet - please wait and try again.

Explanation

This error occurs during the database recovery phase of Adaptive Server startup when you try to access a database for which recovery has not yet been completed. It can also occur if the database was marked suspect due to an earlier problem.

If the database has been marked suspect by an earlier problem, this is a serious error and must be corrected if you want to access your database again.

Action

The time it takes to recover a database depends on:

- The number and size of transactions that must be committed or rolled back
- The amount of memory

Determine whether the recovery phase of the database has completed:

- Examine your Adaptive Server error log
- Execute the following command twice, 20 seconds apart, to help you determine if the databases are still undergoing recovery.

```
> select * from master..sysprocesses where status ="background"
```

When you receive the query's output, check whether *<physical_io>* is increasing over time.

Note A “background” status does not always indicate that the process is executing a recovery. This status can also indicate that tasks such as a threshold procedure or Replication Agent (for example) are running.

If the database is still being recovered, wait until recovery is complete and then try to access the database.

If you are still unable to access the database and the database has been marked suspect because of an earlier problem, determine what the earlier problem was by looking at the Adaptive Server error log. Follow the instructions in this manual for recovering from that error. If no instructions exist, call Sybase Technical Support.

Additional information

Before calling Sybase Technical Support, have the information listed in “Reporting errors” on page 4 available, including the output from:

- `select * from master..sysusages`
- `select dbid, name, status from sysdatabases`

Version in which this error is raised

All versions

Error 924

Severity

14

Message text

Database '%.*s' is already open and can only have one user at a time.

Explanation

This error occurs when an attempt is made to access a database that has been set to single-user mode and is being accessed by another user. Only one user at a time can access a database that is set to single-user mode.

Action

- 1 Verify that the database being accessed is in single-user mode:


```
1> sp_helpdb <database_name>
2> go
```

If the database is in single-user mode you will see “single user” in the status column along with any other options that are set for the database.

- 2 Execute `sp_who` to determine who is accessing the database. Look for the database name under the `dbname` column and the login name in the `loginame` column.
- 3 Contact the individual and arrange for database access or contact your System Administrator (“sa”). The “sa” or database owner (“dbo”) can set the database to multi-user access when the database is not in use by executing the following from the `master` database:

```
1> sp_dboption <database_name>, single, false
2> go
```

```
1> use <database_name>
2> go
```

```
1> checkpoint
2> go
```

Refer to the most recent version of *Reference Manual: Procedures* for information about `sp_dboption`.

Version in which this error is raised

All versions

Error 925

Severity	19
Message text	Maximum number of used databases for each query has been exceeded. The maximum allowed is %d.
Explanation	<p>Adaptive Server has an eight-level-deep stack of open databases for each process. Every time a database is opened, it gets pushed onto the stack, and when it is closed, it is removed from the stack.</p> <p>There is usually one open database for each process. If you run a query that uses a table in a different database, Adaptive Server opens that database just long enough to use the table and then closes the database.</p>

Sometimes, however, a database must remain open longer. For example, if you execute a stored procedure in a different database than the one you are currently using, Adaptive Server will open that database and keep it open until the stored procedure finishes.

Only eight databases can be open at once. A single database opened multiple times is added to the stack each time it is opened. This can result in error 925, even though only one database has been used. If you have stored procedures that call other procedures in other databases, you can exceed this limit.

Note The use `<database_name>` command closes the database that was previously open, so error 925 would not be generated by eight use `<database_name>` commands in a row in Transact-SQL.

An example of a set of queries that exceeds the eight-level limit follows:

```
1> use db9
2> go

1> create proc pr9
2> as
3> select * from table1
4> go

1> use db8
2> go

1> create proc pr8
2> as
3> exec db9..pr9
4> go

1> use db7
2> go

1> create proc pr7
2> as
3> exec db8..pr8
4> go

1> use db6
2> go

1> create proc pr6
2> as
3> exec db7..pr7
4> go

1> use db5
```

```
2> go

1> create proc pr5
2> as
3> exec db6..pr6
4> go

1> use db4
2> go

1> create proc pr4
2> as
3> exec db5..pr5
4> go

1> use db3
2> go

1> create proc pr3
2> as
3> exec db4..pr4
4> go

1> use db2
2> go

1> create proc pr2
2> as
3> exec db3..pr3
4> go

1> use db1
2> go

1> create proc pr1
2> as
3> exec db2..pr2
4> go
```

Action	Be aware of the eight-level open database limit when writing queries. If you believe error 925 is occurring for a reason other than a query that has more than eight database openings, restarting Adaptive Server should clear the problem.
Version in which this error is raised	All versions

Error 926

Severity	14
Message text	Database '%.*s' cannot be opened. An earlier attempt at recovery marked it 'suspect'. Check the SQL Server errorlog for information as to the cause.
Explanation	<p>This error occurs when you attempt to reference a database that has been marked suspect:</p> <ul style="list-style-type: none">• During start-up of Adaptive Server• By the System Administrator as a result of certain critical errors <p>This is a serious error and must be corrected if you want to access your database again.</p> <hr/> <p>Note This error can be caused by a hardware problem.</p> <hr/>
Action	<p>Error 3414 is a related error that occurs during startup when Adaptive Server cannot complete the recovery of a database. Refer to “Error 3414” on page 416.</p> <p>Since the 926 error is the result of an earlier error or action, the recommended action is to determine what caused the database to be marked suspect.</p> <p>In order to determine the cause, check the Adaptive Server error log for error messages for the database in question and try to eliminate those first by using the troubleshooting procedures in this manual.</p> <p>Depending on why the database was marked suspect, you may choose to remove its suspect flag if you are certain that the critical error which caused the database to be marked suspect has been resolved (for example, if one of the database devices was not available when Adaptive Server was started and you are sure that the device is available now). If you choose to reset the suspect status, refer to “Resetting a database’s “suspect” status” in the chapter “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.</p> <p>If you cannot find any procedures recommended for your specific errors, call Sybase Technical Support for assistance.</p>

If the specified database does not contain important data or if you have a known, clean backup of it, you may choose to drop it first, re-create it, and then load the clean database dump into it. Refer to “Dropping a database when drop database fails” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for information on how to drop a database that has been marked suspect.

Before loading the database dump into the newly created database, make sure that the new database and the dumped database have the same data and log mapping, and the same user segment definitions. Refer to “Error 2558” on page 314 for information about how to do this.

Additional information

Before calling Sybase Technical Support, have the information listed in “Reporting errors” on page 4 available, including the output from:

- `select * from master..sysusages`
- `select dbid, name, status from sysdatabases`

Version in which this error is raised

All versions

Error 930

Severity

14

Message text

Database '%.*s' cannot be opened because either an earlier system termination left LOAD DATABASE incomplete or the database is created with 'for load' option. Load the database or contact a user with System Administrator (SA) role.

Explanation

This error occurs when you attempt to reference a database:

- Whose load database process has been terminated before completion, or
- That was created with the `for load` option and that has not been loaded yet.

Action

Eliminate the 930 error by determining which of the above situations applies to you, then follow the appropriate set of steps below.

Previous load database failed

Since the 930 error is, in this case, the result of an earlier error, the recommended action is to determine what caused the load database process to be terminated before completion.

- 1 Check the Adaptive Server and Backup Server error logs for error messages about the database in question, and try to eliminate those first by using the troubleshooting procedures in this manual. If you cannot find any procedures recommended for your specific errors in this manual, call Sybase Technical Support for assistance.
- 2 Examine your operating system error log to determine if hardware errors may have caused the corruption. If errors of this type occurred, solve the problem that caused them.
- 3 Enter the load database command again.

Incorrect use of *create database for load*

Load the database before you try to reference it.

If you do not have enough space allocated to the specified database and get Error 3105 when you attempt to load the database, use one of the following methods to increase the size of the database to be loaded:

- Drop the newly created database first and then re-create it with the size specified in the error 3105 message. If you cannot drop the database with `drop database` because the database has been marked suspect, and you are getting error 926, the instructions in “Dropping a database when drop database fails” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide.

or:

- Use the command `alter database db_name for load`.

To avoid 2558 errors or a remapped database, the database being loaded into must be identical to the dumped database in terms of data, log, and segment mapping. Refer to “Error 2558” on page 314 for more information.

If the specified database is the default database for a certain user, that user will get errors 930 and 4001 and will be put in the `master` database when they connect to the Adaptive Server. Use the following command if you want to change that user’s default database:

```
1> sp_modifylogin <user_name>, defdb, <database_name>
2> go
```

Additional information

Before calling Sybase Technical Support, have the information listed in “Reporting errors” on page 4 available, including the output from:

- `select * from master..sysusages`
- `select * from master..sysdatabases`

Version in which this error is raised All versions

Error 935

Severity 10

Message text WARNING - the timestamp in database '%.*s' is approaching the maximum allowed.

Explanation Each database has one global timestamp which is kept in memory in a structure called a `dbtable`. The timestamp is not a date and time; rather, it is a sequence number which keeps track of modifications of pages within the database. Each time a page in the database is modified, the page gets the current timestamp and the global timestamp gets incremented. The global timestamp “travels” with a database through dumps, loads and upgrades. Timestamps are critical during recovery, to ensure that transactions are properly rolled forward or rolled back. The global timestamp is also used for timestamp columns.

The maximum value for timestamps is very high (`0xffff 0xffffffff`), so it is very unlikely that under normal circumstances a database will reach this limit. However, use of some unsupported database maintenance commands can cause a higher than expected global timestamp value.

Each time `dbtable` is created for the database, Adaptive Server checks whether the timestamp value is approaching the maximum by comparing the current timestamp value with a threshold value defined by Adaptive Server. Error 935 is raised if the timestamp value exceeds this threshold.

Action Error 935 gives you advance warning about a serious impending problem.

Warning! After error 935 is raised, only a few more transactions are possible in the database. If you do not address the problem, the timestamp value will reach the maximum and the database will no longer be usable.

A database dump and load is not an option to address this problem. When you receive error 935, re-create the database in question as soon as possible using these steps:

- 1 bcp out all table data.
- 2 defncopy out all procedures, rules, defaults, and triggers.

- 3 Use the drop database command to drop the database.
- 4 Use the create database command to re-create the database.
- 5 Create all tables.
- 6 bcp in all table data.
- 7 Create indexes.
- 8 defncopy in all procedures, rules, defaults, and triggers.
- 9 Check that the user IDs in sysusers in the restored database are consistent with master..syslogins. Mismatched user IDs in databases can lead to permission problems.

Scripts are recommended for reliable re-creation of the database objects.

Additional information

Refer to the glossary for additional information about timestamps and dbtable. Refer to “Developing a Backup and Recovery Plan” in the *System Administration Guide: Volume 2* for complete information about how to safely re-create databases.

Version in which this error is raised

All versions

Error 940

Severity

26

Message text

Dbtable in wrong state for operation: %S_DBT.

Explanation

Adaptive Server tracks the status of each open database using a memory structure called a dbtable, which contains state information such as the number of open transactions in the database, the number of open objects, keep counts that ensure proper use of each dbtable, and so on. Error 940 is raised when a database operation is attempted which is incompatible with the current state of its dbtable.

Error 940 is due to an internal Adaptive Server problem. It results in a stack trace and prevents further work in the server.

The error is raised with the following states:

State	Meaning
1	Database is being dropped but the dbtable is not marked for dropping.

State	Meaning
2	dbtable was marked for dropping by a task other than the one dropping the database.
3	Database is being dropped but the dbtable's keep count is not zero. This state is only raised by diagserver.
4	Database is being dropped but is not marked active.
5	dbtable is being locked but the database is not active.
6	While decrementing the dbtable keep count, it was found to be less than zero.
7	Illegal attempt by a process to increment the keep count.
8	Similar to state 7.
10	Invalid pointer to dbtable structure.
11	While terminating use of a dbtable, the keep count was found to be less than zero.
12	Attempted to unlock a dbtable which is not currently locked.
13	Attempted to cleanup a dbtable not marked for cleanup.
14	A non-system task attempted to update the keep count.
15	Attempted to update a newly created dbtable that is not marked as created. (Raised by diagserver.)
16	Similar to state 14.
17	Attempted to free a dbtable not marked for cleanup. (Raised by diagserver.)
18	Attempted to free a dbtable whose keep count is not 1.
19	Retrieved a dbtable from the free list, but the dbtable is not marked as free.
20	Illegal attempt to mark a dbtable active (raised by diagserver).
21	Illegal attempt by a process that does not own a dbtable to keep it.
22	Attempted to uninstall a dbtable that is not in the appropriate state (raised by diagserver).
23	Attempted to uninstall a dbtable that is not on the active chain.
25	Database is not in 'drop' state as expected.
26	dbtable was found to be free although it is not marked free.
27	dbtable was found to be free although its keep count is non-zero.
29	dbtable is found to be active although it is not marked active.
30	More than one dbtable is found for a database.

Action

Shut down and restart the server. If the problem persists, call Sybase Technical Support. You may need to upgrade to a version in which the problem is resolved.

Version in which this error is raised

Version 11.5.1 and later

Error 941

Severity

20

Message text

Illegal database context operation.

Explanation

The Adaptive Server Database Context Manager controls access to databases. Any process opening or closing a database must do so through this subsystem, which maintains a list of open databases, the server user ID (suid) accessing the database, and other information in memory. The manager also accesses the process status structure (PSS), which is a memory structure associated with each process describing the state of the process and usage of various resources used by the process.

Error 941 is raised when the Database Context Manager encounters a problem in fulfilling a database context change request.

Error 941 occurs with the following states:

State	Meaning
1	A request is made to close the current database and return the context to the previous database for the requestor. However there is no previous database on the context list.
2	A request is made to close the last database open for the requestor. However the context list indicates that this is not the last open database.
3	A database close request does not include a context for the action (such as returning to the previous database).
4	A request is made to initialize the context list. However the list is not empty and cannot be initialized.

Error 941 is due to corruption of the PSS or other memory structures.

Action

Shut down and restart Adaptive Server to clear any memory corruption. If the problem persists, increase the value of the `number of open databases` parameter.

Version in which this error is raised

All versions

Error 945

Severity	21
Message text	Unable to locate page %ld for database id %d in the in-memory copy of Sysusages (spid=%ld). This is an internal error. Please contact Sybase Technical Support.

Note This error can be caused by a hardware problem.

Explanation When the Adaptive Server Open Database Manager requests a page (for allocating extents and a number of other operations), it receives a page number and a pointer to a disk map array. It then determines the index of the map element that contains the target page and returns that index. When Adaptive Server cannot find the requested page number in the database disk map array, error 945 occurs.

Possible causes of error 945 are:

- Memory corruption.
- Bad allocation page pointers (sysusages table corruption).

Action

- 1 Shut down and restart Adaptive Server. If the corruption is in memory, this should clear the 945 error. If error 945 persists, go on to step 2.
- 2 Check your record of sysusages information for the database in the error message against what is currently in master..sysusages. (If you do not have a record of your sysusages information, try:

```
sp_help <database_name>
```

to get a subset of the sysusages information.) To determine the sysusages information currently held by Adaptive Server, type:

```
1> select * from sysusages where dbid = <database_ID>
2> go
```

where <database_ID> is the number from the error message output.

- 3 If the information you have on record for that database does not match the information currently in sysusages, refer to “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for information about restoring a corrupted sysusages table.
- 4 If the information you have on record for that database matches what is currently in sysusages, the database specified in the error message is corrupt. Restore the database from known, clean backups.

Version in which this error is raised All versions

Error 950

Severity	14
Message text	Database '%.*s' is currently offline. Please wait and try your command again later.
Explanation	<p>The Adaptive Server Database Context Manager controls access to databases. Any process opening or closing a database must do so through this subsystem. Error 950 is raised when the Database Context Manager is unable to honor a database access request because the database is offline, and is not available for general use.</p> <p>During a load database, Adaptive Server takes the database being loaded offline and the database remains in this state until it is brought back online.</p> <p>During an Adaptive Server installation upgrade, a database is marked offline when loading the transaction log that is to be upgraded. You may see 950 errors:</p> <ul style="list-style-type: none">• If the load completed and the upgrade was successful, but the database was not brought back online.• If the upgrade failed, leaving the database offline.
Action	<p>If a load is in progress, complete the load sequence (including any loads of transaction log dumps), then use <code>online database</code> to make the database available for use.</p> <p>If this error is raised due to an upgrade failure, the corrective action depends on several factors:</p> <ul style="list-style-type: none">• The original Adaptive Server version, and the version to which you are upgrading.• The phase of the upgrade process where failure occurred.• Whether the databases being upgraded has tables or other objects marked for replication. If the transaction log is not drained because log records are waiting to be processed by the Log Transfer Manager, the upgrade process for that database will fail, leaving the database offline.

	Refer to “Recovering from a failed installation” in the “Troubleshooting” chapter of the installation and configuration guide for your platform to determine how to proceed.
Additional information	If you use Replication Server, you must disable replication before attempting to upgrade. Refer to your platform’s Adaptive Server installation guide for further information. For example, if you are upgrading, see “Appendix B: Upgrading Servers with Replicated Databases” in the appropriate installation guide..
	When you call Sybase Technical Support, have the information listed in “Reporting errors” on page 4 on hand.
Version in which this error is raised	All versions

Error 965

Severity	21
Message text	Unable to create the database disk map for the archive database '%d'.
Explanation	Before loading an archive database. Adaptive Server scans the <code>sysaltusages</code> table to build a disk map that represents the database to be loaded. A problem was detected when Adaptive Server was attempting to access the table.
Action	Check for any problems with the scratch database or the <code>sysaltusages</code> table in the scratch database. Check for a more precise error regarding the inability to open, lock or read the <code>sysaltusages</code> table.
Additional information	Set trace flags 15701 and 15706 on to provide additional information about remapping pages and <code>sysaltusages</code> . If this error persists, setting up a configurable shared memory dump (CSMD) may be helpful in determining the root cause for this error. Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of <i>Troubleshooting and Disaster Recovery</i> .
Version in which this error is raised	Version 12.5.4 and later

Abstract Plan Manager Errors (1000s)

This chapter covers errors raised when using abstract plans.

An abstract plan describes the execution plan for a query using a language created for that purpose. Adaptive Server can generate an abstract plan for a query, and save the text and its associated abstract plan in the `sysqueryplans` system table. Using a rapid hashing method, incoming SQL queries can be compared to saved query text, and if a match is found, the corresponding saved abstract plan is used to execute the query.

Error	Page
Error 1024	171

Error 1024

Severity	20
Message text	Abstract Plan (AP) : Fatal error in the lexer while scanning the Abstract Plan. Please contact Sybase Technical Support.

Note “lexer” is another term for parser.

Explanation	While scanning the abstract plan, Adaptive Server encountered a fatal error.
Action	<ul style="list-style-type: none"> Examine the abstract plan and correct any syntax errors, then try reusing the plan. If there are no syntax errors, set trace flag 116 to print out the text of the abstract plan that caused this error and contact Sybase Technical Support.

Pre-version 15.0.2:

```
dbcc traceon (116)
execute query
dbcc traceoff(116)
```

Version 15.0.2 and later:

```
set switch on 116 with override
execute query
set switch off 116
```

Additional information

Please have the abstract plan text and the original query available if you contact Sybase Technical Support.

See the most recent version of *Performance and Tuning Series: Query Processing and Abstract Plans* for more information about using abstract plans.

Version in which this error is raised

Version 12.5.3 and later

This section contains error messages for the Adaptive Server Page Manager.

Error	Page
Error 1105	173
Error 1108	183
Error 1120	184
Error 1124	185
Error 1127	186
Error 1129	188
Error 1131	190
Error 1133	191
Error 1141	193
Error 1142	195
Error 1143	199
Error 1159	200
Error 1161	200

Error 1105

Severity

17

Message text

Can't allocate space for object '%.*s' in database '%.*s' because '%.*s' segment is full/has no free extents. If you ran out of space in syslogs, dump the transaction log. Otherwise, use ALTER DATABASE or sp_extendsegment to increase size of the segment.

Explanation

This error occurs when Adaptive Server is unable to allocate space in a database. The state of the error indicates what type of space could not be allocated:

State	Type of space that cannot be allocated
1	Data space

State	Type of space that cannot be allocated
2	Extent allocation space
3	Log space
4	Log space. A job has been terminated because the abort tran on log full database option is set.
5	Log space during an upgrade
6	Log space after the last chance threshold has already been hit
7	Log space
8	Log space after the last chance threshold has already been hit (written to the error log)
9	Log space after the last chance threshold has already been hit (sent to client)
10	Data space – reservation failed
13	Data space - no space found during object allocation map (OAM) hint scan
14	Data space - no space found during full OAM scan
15	Data space - no space found from GAM
16	Data space - no space left in the object's segment (sent to client)
17	Data space - no space left in the object's segment (written to the error log)

1105 errors are written to the error log only when they occur during the checkpoint process. When a user process encounters this error, the message is returned to the client application without being written to Adaptive Server's error log.

To correct the error, you must know:

- The Adaptive Server version level.
- Whether the error occurred during run time or during recovery:
 - Run time is a period during which Adaptive Server performs normal daily activities, such as receiving a query from a client application, processing a transaction, or returning results from a transaction.
 - An Adaptive Server is in recovery when doing activities such as restarting, performing the processing that occurs after a database dump is loaded, and performing the processing that occurs after a transaction dump is loaded.
- The state of the error.

Action

Use the sections below to correct the error, depending on whether the error occurred at run time or during recovery.

Runtime 1105 errors

If the error occurred during runtime, use the next section to correct the error. Refer to “Recovering from 1105 errors” if the error occurred during recovery.

Runtime 1105 errors: States 1, 2, 10, 13, 14, 15, 16, 17

The data segment is full on the indicated data base. To obtain more data space, do one or both of the following:

- Use `alter database` or `sp_extendsegment` to increase the size of the data segment.
- Drop objects from the database or delete rows from a table.

Runtime 1105 errors: State 3, 7

The log segment is full on the indicated database. To clear space in the log, perform these steps:

- 1 Determine approximately how many pages the transaction log occupies using the commands:

```
1> use <database_name>
2> go
```

If the version is 12.5.x:

```
1> select data_pgs (8, doampg)
2> from sysindexes where id=8
3> go
```

If the version is 15.0.x:

```
1> select data_pages(<database id>,8,0)
2> go
```

Where “8” is the ID for `syslogs`. The result is the number of data pages (2K pages on most operating systems) that the transaction log occupies.

Note The query results may be inaccurate by as many as 16 pages, but using the query is much faster than counting the rows in the transaction log.

- 2 Dump the inactive portion of the transaction log using the `dump transaction` command. If this command fails with the 1105 error, retry the command using the `with truncate_only` option. If this command also fails with the 1105 error, run `dump transaction with no_log`.

Note Using `dump transaction with no_log` can result in an 813 error. Therefore, do not use `dump transaction with no_log` unless it is absolutely necessary (when `dump transaction with truncate_only` will not truncate the log).

- 3 Repeat step 1. If the transaction log occupies significantly fewer pages, continue with step 4.

If, however, there are still a large number of pages in the `syslogs` table, an outstanding transaction is probably preventing the log from being cleared. If this is the case, choose one of these options:

- Refer to the section “Detecting and clearing long-running transactions” in Chapter 4 in the most recent version of the *Troubleshooting and Disaster Recovery* guide.
 - Restart Adaptive Server and repeat step 2. When Adaptive Server starts and this database is recovered, the outstanding transaction is rolled back, allowing the log to be cleared by a subsequent `dump transaction`. Depending on the size of the outstanding transaction, rolling back the transaction during startup recovery may be significant.
- 4 If the `dump transaction` command was executed with either the `no_log` or the `truncate_only` options in step 2, dump the database now, because these options prevent subsequent changes recorded in the log from being used to recover from a media failure. You must do a `dump database` now to ensure the recoverability of subsequent database modifications using transaction logs.

Note This database dump is not required if your backup and recovery procedures do not include saving the transaction logs for media failure recovery.

Do not assume that the occurrence of 1105 errors automatically means your transaction log is too small. If the data and the log are on the same segment, either the state 1 and state 2, or the state 3 actions may free enough space without increasing the size of the transaction log.

If you are concerned that your transaction log is not large enough for your Adaptive Server, refer to “Creating and Managing User Databases” in the *System Administration Guide: Volume 2* and read the section “Estimating the Transaction Log Size” before deciding to increase the log size. If appropriate, refer to *Reference Manual: Commands* for more information about the `alter database` command used to increase log size. Remember that once they have been increased, log and database sizes cannot be decreased.

The above procedures may not be effective in resolving the 1105 error due to excessive data growth in mixed data and log segments; you may need to move the data to a different device in this situation.

Runtime 1105 errors: State 4, 6, 8, 9

The log space is full on the indicated database. LCT (last chance threshold) was trying to empty the log but ran out of space trying to log information. Perform the same procedure as in “Runtime 1105 errors: State 3, 7” on page 175.

If `dbcc checktable on syslogs` shows space available in the transaction log, check your `sysusages` system table. Error 1105 may occur if the `unreservedpgs` column for this database has NULL values in it when it should have numeric values. Update that column to 0 (instead of NULL).

Runtime 1105 errors: Replicated databases

You may see 1105 errors when the log is full (states 3 and 4) on a database that is replicated or has been loaded from a replicated database. This occurs because Adaptive Server is unable to truncate the log past the LTM truncation point, even with `truncate log on checkpoint set`.

If you are no longer replicating this database, or are certain that you can resynchronize the replicated database, instruct the server to ignore the LTM truncation point for the database:

```
1> dbcc settrunc (ltm, ignore)
2> go

1> dump transaction with truncate_only
2> go
```

The server can now truncate the log.

Runtime 1105 errors: *tempdb*

1105 errors occur on `tempdb` while Adaptive Server is processing transactions, never during recovery.

The tempdb transaction log behaves just like a user database transaction log with the sp_configure option trunc log on chkpt set to “on”. But data management is more difficult with tempdb because it is more difficult to get information about space use: work tables cannot be referenced by the user, and the log is truncated approximately every 60 seconds. The default size of tempdb is 2MB when Adaptive Server is installed.

Note Read all relevant sections about transaction log management in the “Auditing” chapter of *System Administration Guide: Volume 1* before making any changes at your site.

Possible causes of the 1105 error in tempdb, and their symptoms and solutions include:

Symptom	Possible cause	Solution
Too many transactions fill up tempdb transaction log	Doing too much work in tempdb (or temporary tables) fills up the log.	Design your application for more efficient use of tempdb.
Too many transactions fill up tempdb transaction log	Processing more transactions per 60 seconds than can be logged in tempdb with current space.	Increase size of the tempdb transaction log (refer to alter database in <i>Reference Manual: Commands</i> for details) or perform manual dump transaction with truncate only commands on tempdb.
One transaction is too big	Your transaction may create a Cartesian product that fills the log. There may be too many steps between the begin transaction and the commit transaction. A single statement such as insert, delete or update may involve so many rows that the transaction log fills up before the statement can finish.	Rewrite the transaction to avoid Cartesian products and long individual queries. See the section “Setting up transaction log management” in the Auditing chapter of <i>System Administration Guide: Volume 1 and Performance and Tuning: Monitoring and Analyzing</i> for information and suggested strategies.

Recovering from 1105 errors

If the 1105 error occurred during run time, use the section “Runtime 1105 errors” on page 175 to correct the error. If the error occurred during recovery, use the procedure below which corresponds to the database on which the 1105 error occurred:

- A user database
- The master database

- The model database

Recovery for 1105 errors on a user database

Use the following procedure to correct 1105 errors on a user database which occur during recovery:

- 1 Manually set the database's status to no chkpt on recovery and single-user:

Warning! Because the database was marked suspect on the original recovery attempt, this procedure also resets some internal status bits to allow the database to recover normally. Therefore, do not use this procedure under any other circumstances to avoid additional serious problems.

```

1> sp_configure "allow updates", 1
2> go

1> begin transaction
2> go

1> update master..sysdatabases
2> set status = status & ~256
3> where name = "<database_name>"
4> go

1> update master..sysdatabases
2> set status = status | 4112
3> where name = "<database_name>"
4> go

```

- 2 Check that each of the above update commands affected only one row. If more than one row was affected, issue a rollback transaction. Otherwise, commit the transaction and shut down Adaptive Server:

```

1> commit transaction
2> go

```

```

1> shutdown
2> go

```

- 3 Restart Adaptive Server.
- 4 Dump the transaction log with the `no_log` option and reset the database status:

```

1> use master
2> go

```

```
1> dump tran <database_name> with no_log
2> go

1> sp_dboption <database_name>, "no chkpt", false
2> go

1> sp_dboption <database_name>, single, false
2> go

1> use <database_name>
2> go

1> checkpoint
2> go

1> sp_configure "allow updates", 0
2> go
```

Recovering from 1105 errors on the *master* database

If an 1105 error occurs on the *master* database during recovery and Adaptive Server does not run, correct the problem with the following procedure:

- 1 Create a *runserver* file that includes the 3607 trace flag. This trace flag allows Adaptive Server to restart without going through normal recovery, and should only be used in this procedure. Refer to “Starting Adaptive Server with trace flags” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for instructions.
- 2 Start Adaptive Server with the 3607 trace flag using the *runserver* file created in step 1. The recovery process performs the minimum actions required to allow access to each database, including *master*.
- 3 Once Adaptive Server has restarted and the *master* database is accessible, dump the transaction log with the *no_log* option and shut down Adaptive Server:

```
1> dump tran master with no_log
2> go
```

```
1> shutdown with nowait
```



```
2> go
```

Note It is mandatory that you use `shutdown` with `nowait` here. Failure to use the `with nowait` option will suicide every database's transaction log.

- 4 Restart Adaptive Server with the original `runserver` file. The database should recover normally. Now dump the master database:

```
1> dump database master
2> to <logical_dump_device>
3> go
```

Recovering from 1105 errors on the *model* database

If an 1105 error occurs on the `model` database during recovery, Adaptive Server may not start. The error occurs because the `tempdb` database is required to start Adaptive Server, but it cannot be built due to the problem with `model`. To correct the problem and restart Adaptive Server follow these steps:

- 1 Create a new `runserver` file that includes the 3608 trace flag. This trace flag causes Adaptive Server to recover only the master database. Refer to the section “Starting Adaptive Server with trace flags” in the chapter “Other Useful Tasks” of the most recent version of *Troubleshooting and Disaster Recovery* for instructions.
- 2 Start Adaptive Server with the 3608 trace flag using the `runserver` file created in step 1.
- 3 Manually set model status to `no chkpt on recovery and single-user`:

```
1> sp_configure "allow updates", 1
2> go
```

```
1> begin transaction
2> go
```

```
1> update master..sysdatabases
2> set status = status & ~256
3> where name = "model"
4> go
```

```
1> update master..sysdatabases
2> set status = status | 4112
3> where name = "model"
4> go
```

- 4 Verify that each of the above update commands affected only one row. If more than one row was affected, issue a rollback transaction. Otherwise, commit the transaction and shut down Adaptive Server:

```
1> commit transaction
2> go
```

```
1> shutdown with nowait
2> go
```

- 5 Restart Adaptive Server, using your regular *runserver* file.
- 6 Execute the following command to truncate the inactive portion of the transaction log in model:

```
1> dump tran model with no_log
2> go
```

- 7 Reset the database options on model:

```
1> sp_dboption model, "no chkpt", false
2> go
```

```
1> sp_dboption model, single, false
2> go
```

```
1> use model
2> go
```

```
1> checkpoint
2> go
```

- 8 Disable updates to the system catalog and shut down Adaptive Server:

Warning! It is mandatory that you use `shutdown with nowait` here. Failure to use the `with nowait` option will suicide every database's transaction log.

```
1> sp_configure "allow updates", 0
2> go
```

```
1> shutdown with nowait
```

```
2> go
```

9 Restart Adaptive Server.

Version in which this error is raised

All versions

Error 1108

Severity

21

Message text

Cannot deallocate extent %ld, database %d. Object id %ld, index id %d, status %d in extent does not match object id %ld, index id %d, status %d in object being deallocated. Run DBCC CHECKALLOC.

Explanation

Error 1108 occurs when Adaptive Server is deallocating an extent, if the extent information on the allocation page does not match the information for the object being deallocated. This can be caused by corruption of the allocation page or memory corruption.

A software problem can cause this error to occur:

- During rollback of a transaction
- When rows for an object are being deleted or updated
- When an object is being dropped

This is a serious error as it indicates corruption has occurred.

Note Prior to version 12.5, error 1108 was only raised by diagserver.

Action

- 1 Run `dbcc checkalloc` on the database named in the message. Resolve any errors reported by `dbcc checkalloc` using information about those errors in this manual.
- 2 If errors still exist, restore the database from a clean backup or call Sybase Technical Support.

Additional information

If you need assistance from Sybase Technical Support, have the information listed in “Reporting errors” on page 4 available when you call.

Version in which this error is raised

All versions

Error 1120

Severity	20
Message text	Attempt to call pg_readalloc on page %ld which is not an allocation page.
Explanation	<p>Allocation pages are used to track space in the database. Each allocation page for an object has an entry in the object allocation map (OAM) for that table or index. The OAM entry for each allocation page stores the number of used and unused data or index pages on that allocation page.</p> <p>Error 1120 occurs when you have an invalid page number for an allocation page in the allocation section of an OAM. In this case, Adaptive Server does not have accurate information about which pages have been allocated.</p> <hr/> <p>Warning! This is a serious error because Adaptive Server cannot correctly insert or update data when it occurs. Action should be taken immediately or further database corruption will occur.</p> <hr/>
Action	<p>Use the page number in the error message to identify the object and the database that is affected by the error. Use the procedure “How to find an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of <i>Troubleshooting and Disaster Recovery</i>.</p> <hr/> <p>Warning! The procedure referenced in the preceding paragraph uses the <code>dbcc page</code> command. Use the <code>dbcc page</code> command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.</p> <hr/> <p>To recover from error 1120, restore your database either using <code>bcp</code> or from clean backups. Refer to “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of <i>Troubleshooting and Disaster Recovery</i>.</p> <hr/>
Additional information	For information about OAM pages, refer to “Checking Database Consistency” in the <i>System Administration Guide: Volume 2</i> , and see the section “Understanding page and object allocation concepts.”

If you need assistance from Sybase Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the output of dbcc page for the corrupted page and the output of dbcc checkalloc for the database identified by dbcc page.

Version in which this error is raised

All versions

Error 1124

Severity

20

Message text

Failed to get lock on allocation page %ld.

Explanation

Adaptive Server uses allocation pages to track space in the database. Allocation pages are updated when an object’s space requirements change.

Adaptive Server raises error 1124 due to the following sequence of events:

- A server process acquires a lock on an object’s allocation page to allocate pages for the object.
- An exception occurs during allocation that requires “undoing” the allocation steps.
- The server is unable to acquire a lock necessary to abort the allocation, resulting in error 1124.

This error may be accompanied by a kernel error (“Current process infected with %d”); error 945 (“Unable to locate page %ld for database id %d in the in-memory copy of sysusages (spid=%ld). This is an internal error. Please contact Sybase Technical Support.”); error 1205 (“Your server command (family id #%d, process id #%d) was deadlocked with another process and has been chosen as deadlock victim. Re-run your command.”) or other errors.

Error 1124 is raised when Adaptive Server has a problem rolling back space allocation.

Action

When error 1124 is accompanied by a 945 error on tempdb (dbid = 2) that is not resolved by shutting down and restarting the server, there may be insufficient space on tempdb or a problem with tempdb storage.

	Resolve these problems using the instructions in the section “Reducing the size of tempdb” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
	In all other instances of the 1124 error, contact Sybase Technical Support.
Additional information	Have the information listed in “Reporting errors” on page 4 ready before calling Sybase Technical Support:
Version in which this error is raised	Version 11.5.x and earlier

Error 1127

Severity	22
Message text	Internal Error: The OAM has no room for new inserts. Check OAM for object %ld, index %d.
Explanation	<p>The object allocation map (OAM) tracks information about the allocation and deallocation of storage space for tables and indexes. At least one OAM page exists for each table (except <code>syslogs</code> and <code>sysgams</code>) and for the table’s indexes. The OAM page has an entry for each allocation page used by the object. The first OAM page for an object has a maximum of 240 entries, and any subsequent OAM pages have a maximum of 250 entries per page.</p> <p>When a new page is allocated for the table or index:</p> <ul style="list-style-type: none">• The appropriate OAM page is updated.• If necessary, a new OAM entry is created for the allocation page which controls the page being added.• If necessary, a new OAM page is created and linked in to the current OAM page chain for the object. This occurs when there are no available slots for a new OAM entry in any existing OAM page for the object.• The total number of entries in the object’s OAM (corresponding to the total number of allocation pages with extents for this object) is adjusted during table updates and also during database recovery. <p>Error 1127 is raised:</p> <ul style="list-style-type: none">• If the total number of entries for all OAM pages for the object exceeds its upper limit (overflow in OAM page header). This can happen if the table grows and occupies space on many allocation pages.

Action

- If the total number of entries is corrupted. This is due to an Adaptive Server problem.

There are different ways to correct this error, depending on the conditions under which it occurred and the nature of the OAM problem. Other errors may also be raised at the same time.

- If this error is raised when you attempt to start Adaptive Server and the Server hangs, the database may be marked suspect; contact Sybase Technical Support for assistance.
- If the error occurs during a table update (for example, during a bulk copy operation), follow these steps:

- a If the object is a system table (a system table's object ID is less than 100) and the index ID is not 0 (zero), refer to "Fixing a corrupted system table index" in the most recent version of *Troubleshooting and Disaster Recovery* for instructions on how to repair the system table index.

If the object is a system table and the index ID is 0 (zero), contact Sybase Technical Support. They may be able to help you repair the corruption but it may be necessary to restore the database from clean backups.

If the object encountering the error is not a system table, continue to step 2.

- b If the error occurs on a user table, the problem may be an overflow in the OAM page header or corruption of the page header. Run the following command to display OAM page information.

Note You must have the `sybase_ts_role` granted to you to run this command and run `dbcc traceon(3604)` first.

```
1> dbcc traceon(3604)
2> go
```

(For version 15.0.2 and later, on line 1, run `set switch on 3604` instead.)

```
1> dbcc listoam(database_name, object_ID, index_ID)
2> go
```

For example:

```

1> dbcc listoam(pubs2, 104001776, 0)
2> go
-----
Objid:   104001776      indid:         0
OAM pg cnt:      1      Entry cnt:      1
Rows:           18      Rows Per pg:    6
Used pgs:        4      Unused pgs:     4
OAM status bits set:  PG_OAMPG, PG_OAMSORT
OAM pg #         520 has the following entries
                    (allocpg:used/unused):
                    512: 4/ 4
-----
DBCC execution completed. If DBCC printed error
messages, contact a user with System
Administrator
(SA) role.

```

Call Sybase Technical Support with this information.

Additional information

For information about OAM pages, refer to “Checking Database Consistency” in the *System Administration Guide: Volume 2*, and see the section “Understanding page and object allocation concepts.”

Before calling Sybase Technical Support, have the information on hand that is listed in “Reporting errors” on page 4, including the output of dbcc listoam for the table.

Version in which this error is raised

All versions

Error 1129

Severity

10

Message text

The requested update to the OAM for object %ld, dbid %ld, index %d, oam page %ld, would cause a negative page count. allocation page: %ld, oamarray used: %ld, oamarray unused: %ld, deltaused: %ld, deltaunused: %ld, spid: %ld. Contact your System Admin.

Explanation	<p>Most tables and all indexes have at least one object allocation map (OAM) page. One OAM page holds information about up to 250 allocation pages (248 if the object is an identity column). Each allocation page for an object has an entry in the OAM pages for that table or index. The OAM entry for each allocation page stores the number of used and unused data and index pages on that allocation page.</p> <p>When a new page is needed or a page is deallocated, the OAM entry has to be updated. The number of used and unused pages on the allocation page is not allowed to be less than 0. Error 1129 occurs if the number becomes less than 0.</p>
Action	<ol style="list-style-type: none"> 1 Look at the error message to determine the object ID and the index ID. <ul style="list-style-type: none"> If the object that encountered the error is not a system table (a system table's object ID is less than 100), go to step 2. If the object is a system table and the index ID is not 0 (zero), refer to "Fixing a corrupted system table index" in the chapter "Other Useful Tasks" in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide for instructions on how to repair the system table index. If the index ID is 0 (zero), contact Sybase Technical Support. They may be able to help you repair the corruption, but it may be necessary to restore the database from clean backups. 2 For user tables, if the index ID is 0 or 255, go to step 3. <ul style="list-style-type: none"> If the index ID is not 0 or 255: <ol style="list-style-type: none"> a Translate it into an index name: <pre style="margin-left: 40px;">1> use <database_name> 2> go 1> select name from sysindexes 2> where id = <object_ID> and indid = <index_ID> 3> go</pre> b To ensure that the information needed to re-create the index is available, run <code>sp_helpindex</code> on the index prior to dropping it. c Drop the index. d Re-create the index. This clears the corruption in most cases. e Run <code>dbcc checktable</code> on the table to verify that the corruption is gone. 3 If the index ID is 0 or 255, do one of the following: <ul style="list-style-type: none"> • Restore the database from clean backups.

- Refer to “Rescuing data from a corrupted table” in the most recent version of the “Other Useful Tasks” chapter in the *Troubleshooting and Disaster Recovery* guide.

Additional information For more information about OAM pages, refer to the section “Understanding Page and Object Allocation Concepts” within the chapter “Checking Database Consistency” in the *System Administration Guide: Volume 2*.

If you need to call Sybase Technical Support, have the information ready that is listed in “Reporting errors” on page 4, including the output from dbcc traceon (for version 15.0.2 and later, use “set switch on 3604” instead of dbcc traceon.), and dbcc listoam. You must have the “sybase_ts_role” granted to you to run dbcc listoam.

For example, enter the following for pre-15.0.2 versions:

```
1> dbcc traceon(3604)
2) go
```

```
1) dbcc listoam(<database_name>, <object_ID>, <index_ID>)
2) go
```

Warning! Use the dbcc listoam command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

Version in which this error is raised All versions

Error 1131

Severity 22

Message text The OAM for object %ld, index %d, oam page %ld, database=%d, is currently allocated to another object.

Explanation	<p>The object allocation map (OAM) tracks information about the allocation and deallocation of storage space for tables and indexes. At least one OAM page exists for each table (except <code>syslogs</code> and <code>sysgams</code>) and for the table's indexes. Each allocation page used by the object has an entry in the OAM page, showing the number of used and unused data/index pages on that allocation page.</p> <p>Error 1131 is raised when Adaptive Server attempts to validate an OAM page belonging to an object, and determines that the page is actually allocated to a different object.</p> <p>Error 1131 occurs with the following states:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left; padding: 5px;">State</th> <th style="text-align: left; padding: 5px;">Meaning</th> </tr> </thead> <tbody> <tr> <td style="text-align: center; padding: 5px;">1</td> <td style="padding: 5px;">While updating the used and unused counts for an OAM entry, the OAM page is found to be allocated to a different object. This error is raised only by <code>diagserver</code>.</td> </tr> <tr> <td style="text-align: center; padding: 5px;">2</td> <td style="padding: 5px;">(Version 15.0 and later) While checking the validity of an OAM page, the page is allocated to a different object.</td> </tr> <tr> <td style="text-align: center; padding: 5px;">3</td> <td style="padding: 5px;">(Version pre-15.0) While checking the validity of an OAM page, the page is allocated to a different object.</td> </tr> </tbody> </table>	State	Meaning	1	While updating the used and unused counts for an OAM entry, the OAM page is found to be allocated to a different object. This error is raised only by <code>diagserver</code> .	2	(Version 15.0 and later) While checking the validity of an OAM page, the page is allocated to a different object.	3	(Version pre-15.0) While checking the validity of an OAM page, the page is allocated to a different object.
State	Meaning								
1	While updating the used and unused counts for an OAM entry, the OAM page is found to be allocated to a different object. This error is raised only by <code>diagserver</code> .								
2	(Version 15.0 and later) While checking the validity of an OAM page, the page is allocated to a different object.								
3	(Version pre-15.0) While checking the validity of an OAM page, the page is allocated to a different object.								
Action	<p>1131 errors are often seen when you use the system functions <code>data_pgs</code>, <code>reserved_pgs</code>, <code>used_pgs</code>, or <code>rowcnt</code> in a very active database, or a highly volatile database such as <code>tempdb</code>, where objects are being created and dropped frequently. The error can occur when:</p> <ul style="list-style-type: none"> • You use an invalid <code><object id></code> in the function • The page is allocated to another object, or is uninitialized <p>Use trace flag 1116 to suppress the 1131 error when using these system functions in an active or volatile database environment.</p>								
Additional information	<p>See “Starting Adaptive Server with trace flags” in the “Other Useful Tasks” chapter of the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.</p>								
Version in which this error is raised	<p>All versions</p>								

Error 1133

Severity	22
Message text	Page %ld was expected to be an OAM page for %ld and it

is not.

Explanation

Every object allocation map (OAM) page has a unique number in the database. Error 1133 occurs when Adaptive Server attempts to retrieve an OAM page by specifying the page number, but fails because there is no OAM page in the current database with that page number.

Error 1133 can occur when `checkalloc` is run if what is expected to be the OAM page in the OAM page linkage is not an OAM page. In Adaptive Server Enterprise version 12.5, the linkage is accessed via the `doampg` or `ioampg` column for the object in `sysindexes`. (`doampg` is the `sysindexes` pointer to the data page', OAM. `ioampg` is the `sysindexes` pointer to the index pages' OAM.)

Starting with version 15.0, the linkage is accessed via the `datoampage` or `indoampage` column for the object in `syspartitions`. This error can also occur when, due to a problem with Adaptive Server, the OAM page number that is requested is greater than the maximum page number for the database. In these cases, error 1133 is a serious error and may result in reading the wrong OAM page into data cache and corrupting a buffer in the process.

Action

- 1 Use the procedure "Finding an object name from a page number" in the chapter "Other Useful Tasks" in the most recent version of the *Troubleshooting and Disaster Recovery* guide to translate the page number displayed in the error message into an object ID and an index ID.

Warning! The procedure referenced in step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided "as is" without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 If the corruption is on a user table's index (object ID is 100 or greater, and the index ID is between 1 and 250), dropping and re-creating the index should clear the corruption. Use `sp_helpindex` to list all indexes on the table, then rebuild all the nonclustered indexes using either of these two methods:
 - If the index ID is between 1 and 250, drop and re-create each nonclustered index on the table.

or

- If a clustered index also exists on the table, drop and re-create it, which causes all nonclustered indexes to be automatically rebuilt. If your table is large, you may not have the space to do this (a rule of thumb is that 150 percent of your table size must be available).
- 3 If the corruption is not on an index page, run `dbcc tablealloc` or `dbcc checkalloc` on the corrupted object. The output should report OAM allocation errors. Try to fix the reported allocation errors first by referring to the appropriate write-ups in this manual.

If the 1133 errors persist, call Sybase Technical Support for assistance.

Additional information

Before calling Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the output from `dbcc page` for the corrupted page and the output from `dbcc checkalloc` for the database identified by `dbcc page`.

Version in which this error is raised

All versions

Error 1141

Severity

26

Message text

Unexpected value returned to page manager routine: %ld.

Explanation

Adaptive Server uses object allocation map (OAM) pages to track space allocation for Adaptive Server objects. Error 1141 occurs when Adaptive Server tries to update an OAM page following page allocation or deallocation, but the update fails.

Error 1141 can be raised:

- When the database transaction log is full, so that the necessary log records cannot be written to record page allocations
- Due to an Adaptive Server problem

The error is raised in the following states, which indicate the specific page manager activity where the OAM update failed. Pages in a given state may involve all-pages-locked (APL) tables, data-only-locked (DOL) tables, temporary work tables or sort tables, or combinations thereof.

State	Meaning
1	There was an unexpected error when determining whether to update the page counts.
2	OAM update failed when allocating extents; also when deallocating one or more pages in an extent.
3	Update failed when doing post-commit cleanup following deallocation.
4	Update failed when allocating an extent for an APL table.
6	Update failed when deallocating an extent for an APL table.
7	Update failed when deallocating an extent for an APL table, and there were no unused pages in extent.
8	Update failed when performing OAM reallocation. Applies to both APL and DOL tables.
9	Update failed when deallocating an extent for sort pages.
10	Update failed when committing large-scale page deallocations (for example, extents with no pages allocated).

When error 1141 occurs, a stack trace is written to the error log and the current process is terminated.

Action

Other errors may sometimes accompany the 1141 error. For example, error 1105 (see “Error 1105” on page 173) can also be raised when the tempdb transaction log becomes full. If other errors are reported, correct them using the information about those errors elsewhere in this manual. Then retry the transaction.

To check if error 1141 resulted from a full transaction log, use the procedures outlined in the Chapter 6 section “Estimating the transaction log size” of the *System Administration Guide: Volume 2*.

If the 1141 error persists after you have checked for other errors and corrected any log space problems, call Sybase Technical Support.

Additional information

Have the information ready that is listed in “Reporting errors” on page 4 before calling Technical Support.

Version in which this error is raised

All versions

Error 1142

Severity	22
Message text	Invalid OAM Page %ld. Found pstat=0x%x, object=%ld, database=%d, indid=%d.
Explanation	Depending on on the cause, error 1142 can be either a serious error due to actual corruption in the database or a result of misusing a system function (<code>data_pgs</code> , <code>reserved_pgs</code> , <code>rowcnt</code> , or <code>used_pgs</code>). Error 1142 can occur during run time as well as during recovery (during Adaptive Server startup, load database, or load transaction).

Incorrect system function use

The error can occur from incorrect use of any of the following system functions that retrieve information about allocation pages:

- `data_pgs` (version 12.5.x) or `data_pages` (version 15.0 and later)
- `reserved_pgs` (version 12.5.x) or `reserved_pages` (version 15.0 and later)
- `rowcnt` (version 12.5.x) or `row_count` (version 15.0 and later)
- `used_pgs` (version 12.5.x) or `used_pages` (version 15.0 and later)

For example, if you incorrectly type the following query using version 12.5:

```
1> use pubs2
2> go

1> select reserved_pgs (id, root) from
2> master..sysindexes
3> where id = <object_ID>("<table_name>")
4> go
```

error 1142 is raised because the `reserved_pgs` system function attempted to retrieve an invalid OAM page. (Replacing `root` in the above example with a correct OAM page, for example `doampg`, would give expected results.)

Incorrect OAM page entries

If you get error 1142 and you did not use system functions across databases, this error is a sign of serious corruption in your database. To verify this, run `dbcc tablealloc` on the object or `dbcc checkalloc` on the database as discussed in “Incorrect OAM page entry in `sysindexes`” on page 197, step 3.

Adaptive Server uses OAM pages to track how space is allocated for Adaptive Server objects. Error 1142 occurs when Adaptive Server tries to retrieve an OAM page and the status of that page indicates that it is not an OAM page. It occurs with the following states:

State	Meaning
1	The status of the first OAM page indicated that it was not a valid OAM page.
2	Adaptive Server encountered an invalid OAM page during recovery or during an undo or a redo of a truncate table command.
3	Adaptive Server encountered an invalid OAM page while preparing to read and install an OAM page in the buffer cache.
4	Adaptive Server encountered an invalid OAM page during deallocation of an object page.

Action

Recovery from this error depends on the cause. Follow the instructions below, depending on the error's origin.

Incorrect system function use

When you want to use the system functions specified in “Incorrect system function use” on page 195 across databases, select the database to be accessed with a use database command, then run the query for the current database.

Recovery

If error 1142 occurred during the threshold accounting phase of recovery, the database is marked “suspect” and the object referenced in the error message output is corrupted. To determine whether the error occurred during this phase of recovery, look in the error log to determine whether the error occurred immediately after recovery.

- Messages like the following would appear:

```
00:00000:00001:2007/09/26 12:43:36.15 server Recovering database
'testdb'.
00:00000:00001:2007/09/26 12:43:36.27 server Redo pass of recovery has
processed 6 committed and 2 aborted transactions.
00:00000:00001:2007/09/26 12:43:36.33 server Error: 1142, Severity: 22,
State: 1
00:00000:00001:2007/09/26 12:43:36.33 server Invalid OAM Page 64. Found
pstat=0x1111, object=2, database=0, indid=0.
```

If error 1142 occurs after messages like these, perform the following steps to recover from the error. Otherwise, go to “Incorrect OAM page entry in sysindexes” on page 197.

- 1 Turn off threshold accounting for the database affected by the 1142 error so that it can complete recovery:

```
1> use master
2> go
```

```
1> sp_dboption <database_name>, "no free space acctg", true
2> go
```

Where *<database_name>* is the name of the database in the error message.

- 2 For the affected database, use the instructions in “Resetting a database’s “suspect” status” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* guide to reset the status. Shut down and restart the server to complete recovery.
- 3 The object named in the error message is still corrupted. To recover the data, follow the instructions in “How to rescue data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of the of the *Troubleshooting and Disaster Recovery* guide to recover the data.

Note Be sure to perform step 3 before going to step 4. Otherwise, the next time you restart Adaptive Server, the error will reoccur in the database.

- 4 Turn threshold accounting back on for the database affected by the 1142 error:

```
1> use master
2> go
```

```
1> sp_dboption <database_name>, "no free space acctg", false
2> go
```

```
1> use <database_name>
2> go
```

```
1> checkpoint
2> go
```

Where *<database_name>* is the name of the database in the error message.

Incorrect OAM page entry in *sysindexes*

In this case, error 1142 is a serious error and may result in reading the wrong OAM page into the data cache and corrupting a buffer in the process.

Follow these steps to clear error 1142:

- 1 Use the procedure described in “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide to get dbcc page output for the page specified in the 1142 output.

If the page status bits row has the value 0x8000, the page is a valid OAM page. If the value 0x8000 does not appear in the dbcc page output for the page status bits row, the page is not a valid OAM page and the object is corrupted.

Note If the object is not corrupted, call Sybase Technical Support to determine why the 1142 error occurred.

- 2 If the corruption is on a user table’s index (object ID is 100 or greater and index ID is between 1 and 250), run dbcc checktable to verify that the data page linking is good. If dbcc checktable runs cleanly, dropping and re-creating the index should clear the corruption. Use sp_helpindex to list all indexes on the table, then rebuild all the affected nonclustered indexes using one of two methods: :

- If index ID is between 1 and 250, drop and re-create each nonclustered index on the table.

or:

- If a clustered index also exists on the table, drop and re-create it, which causes all nonclustered indexes to be automatically rebuilt. If your table is large, you may not have the space to do this (a rule of thumb is that 150 percent of your table size must be available).

- 3 If the corruption is not on an index page, run dbcc tablealloc or dbcc checkalloc on the corrupted object. The output should report OAM allocation errors. Try to fix the reported allocation errors by referring to the information in this guide.

Note If these instructions do not correct the error, or the error persists, contact Sybase Technical Support.

Additional information

Before calling Sybase Technical Support, have the information available that is listed in “Reporting errors” on page 4, including output from dbcc page for the corrupted page and dbcc tablealloc or dbcc checkalloc for the corrupted object.

Version in which this error is raised All versions

Error 1143

Severity	22
Message text	Invalid OAM Page %ld. Expected object=%ld. Found object=%ld, database=%d, pstat=0x%x, indid=%d.
Explanation	<p>This error occurs when the status on the page is not for an object allocation map (OAM) page or the page is an OAM page and the object ID on the page does not match the object to which the OAM page supposedly belongs.</p> <p>Error 1143 is a serious error and may result in reading the wrong OAM page into data cache and corrupting a buffer in the process.</p>
Action	<ol style="list-style-type: none"> 1 Make a note of the object ID and index ID in the error message output. 2 If the corruption is on a user table's index (object ID is 100 or greater and index ID is between 1 and 250), run <code>dbcc checktable</code> to verify that the data page linkage is good. If <code>checktable</code> runs cleanly, dropping and re-creating the index should clear the corruption. <ul style="list-style-type: none"> Use <code>sp_helpindex</code> to list all indexes on the table and then rebuild all the affected nonclustered indexes using one of two methods: <ul style="list-style-type: none"> • If index ID is between 1 and 250, drop and re-create each affected nonclustered index on the table. or: <ul style="list-style-type: none"> • If a clustered index also exists on the table, drop and re-create it, which causes all nonclustered indexes to be automatically rebuilt. If your table is large, you may not have the space to do this (a rule of thumb is that 150 percent of your table size must be available). 3 If the corruption is not on an index page, run <code>dbcc tablealloc</code> or <code>dbcc checkalloc</code> on the corrupted object. The output should report OAM allocation errors. Try to fix the reported errors first by referring to the appropriate write-ups in this manual. <p>If the 1143 errors persist, call Sybase Technical Support.</p>

Additional information	Before calling Sybase Technical Support, have the information available that is listed in “Reporting errors” on page 4, including output from dbcc page for the corrupted page and dbcc tablealloc or dbcc checkalloc for the corrupted object.
Version in which this error is raised	All versions

Error 1159

Severity	20
Message text	Attempt to locate LSA extent partition control block failed for db id '%1!', objid '%2!', indid '%3!' and partition id '%4!'.
Explanation	This error can occur when Adaptive Server is doing large scale allocation (LSA) using extent partition control blocks. This error is raised when Adaptive Server fails to locate an extent partition control block for a particular (object,index,partition)
Action	Contact Sybase Technical Support.
Additional information	If the error persists, a configurable shared memory dump (CSMD) may be helpful in determining the root cause. Information about CSMD can be found in the chapter “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 1161

Severity	20
Message text	Invalid OAM entry found at slot %d for table ID %d in database `%.*s` (ID %d). This entry is either out of order or its allocation page number is wrong.

Explanation	<p>This error occurs when a problem is detected for an object allocation map (OAM) page used by a table. The error indicates that either the mapping to allocation pages is scrambled (normally allocation pages in use by the object should be mapped in increasing numeric order) or an allocation page being pointed to has the wrong page number.</p>
Action	<p>To clear the problem, and to check whether there are any additional problems with the table, run <code>dbcc tablealloc</code> (with the <code>fix</code> option) on the table mentioned in the error message. For example, if the error mentions the <code>titles</code> table, run:</p> <pre>dbcc tablealloc(titles,full,fix)</pre> <p>If additional errors are detected, check the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide for appropriate actions to take to resolve the problems, or contact Sybase Technical Support for assistance.</p>
Additional information	<p>Each table and index on a table has an OAM. The OAM is stored on pages allocated to the table or index and is checked when a new page is needed for the index or table. One OAM page can hold allocation mapping for 2,000 to 63,750 data or index pages.</p> <p>The OAM pages point to the allocation page for each allocation unit where the object uses space. The allocation pages, in turn, track the information about extent and page usage within the allocation unit. In other words, if the <code>titles</code> table is stored on extents 24 and 272, the OAM page for the <code>titles</code> table points to pages 0 and 256.</p> <p>Additional details on OAM pages can be found in the chapter “Checking Database Consistency” in the <i>Adaptive Server Enterprise System Administration Guide: Volume 2</i>.</p>
Version in which this error is raised	Version 15.0 and later

Lock Manager Errors (1200s)

This section contains error messages for the Lock Manager. The Adaptive Server Enterprise Lock Manager reports locks, deadlocks, lock promotions, and lock contention. See “Lock Management” in the most recent version of the *Performance and Tuning: Monitoring and Analyzing* guide.

Error	Page
Error 1203	203
Error 1204	204
Error 1205	206
Error 1243	208
Error 1244	208
Error 1249	209
Error 1251	210
Error 1265	211
Error 1279	212

Error 1203

Severity	20
Message text	<pre>Caller of lock manager is incorrectly trying to unlock an unlocked object. spid=%d locktype=%d dbid=%d object id or page number=%ld.</pre>
Explanation	<p>Before the Lock Manager tries to unlock an object, it verifies that the object is locked. When an object is not locked during this check, error 1203 occurs.</p> <p>This error is usually due to an Adaptive Server problem that causes an incorrect session descriptor pointer or logical page number to be passed to the Lock Manager.</p> <p>Error 1203 can occur during an insert, update, or select, particularly on a table with an index.</p>

Action	Call Sybase Technical Support.
Additional information	Before calling Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the output of <code>sp_lock</code> .
Version in which this error is raised	All versions

Error 1204

Severity	17
Message text	SQL Server has run out of LOCKS. Re-run your command when there are fewer active users, or contact a user with System Administrator (SA) role to reconfigure SQL Server with more LOCKS.
Explanation	This error occurs when Adaptive Server runs out of locks. Beginning with Adaptive Server Enterprise version 12.5.2, 1204 errors are written to the error log with additional information, including the client IP address, user name and SQL text.

The number of locks available is controlled by the Adaptive Server configuration parameter `number of locks`. Following is an example of the output from `sp_configure`, and a brief description of the output related to locks:

```
1> sp_configure "number of locks"
2> go
```

Parameter Name	Default	Memory Used	Config Value	Run Value
number of locks	5000	0	5000	5000

- The **Default** column contains the default value for `number of locks`, 5000.
- The **Memory Used** column indicates the amount of memory used by the configured locks. Each lock requires 120 bytes of memory.
- The **Config Value** column contains the value to which the `number of locks` configuration parameter has been set with `sp_configure`. If the value has not been explicitly configured, there is a 0 (zero) in this column and the default value is used.
- The **Run Value** column contains the value Adaptive Server is using.

Action	<p>Rerun the command when there are fewer active users, or ask the Sybase System Administrator (“sa”) to reconfigure Adaptive Server for more locks. Increasing the number of available locks impacts performance and memory requirements.</p> <hr/> <p>Warning! Before making any changes to Adaptive Server, refer to the chapter “Memory Use and Performance” in the <i>Performance and Tuning: Basics</i> and the chapter “Configuring Memory” in the <i>System Administration Guide: Volume 2</i> for information about how Adaptive Server uses memory resources.</p> <hr/> <p>To increase the number of locks available, complete these steps:</p> <ol style="list-style-type: none"> 1 Determine the number of locks currently configured for Adaptive Server: <pre style="margin-left: 40px;">1> use master 2> go 1> sp_configure "number of locks" 2> go</pre> <hr/> <p>Warning! If you increase the number of locks available by too many locks, Adaptive Server may run out of memory and be unable to restart. If you run out of memory and cannot restart Adaptive Server, see the section “Adaptive Server does not start after altering configuration” in the most recent version of <i>Troubleshooting and Disaster Recovery</i>.</p> <hr/> 2 Choose the number of locks you want to configure and issue this command: <pre style="margin-left: 40px;">1> sp_configure "number of locks", <new_value> 2> go</pre> <p>Because the number of locks parameter is dynamic, you do not have to restart Adaptive Server to operate with the new values.</p>
Additional information	<p>Refer to the <i>Performance and Tuning: Locking</i> guide for information about locks and other Adaptive Server resources.</p> <p>Refer to the section “number of locks” in the <i>System Administration Guide: Volume 1</i>.</p>
Version in which this error is raised	All versions

Error 1205

Severity 13

Message text Your server command (family id %#d, process id %#d) encountered a deadlock situation. Please re-run your command.

Explanation This error occurs when a process tries to acquire a lock on an object that is locked by a second process when the second process is waiting for a lock on an object that has been locked by the first process. This situation is a deadlock, and can involve more than two processes.

Adaptive Server detects this situation, rolls back the transaction that has accumulated the least amount of CPU time, and notifies the application program of this action with error 1205. This allows other users' processes to move forward.

Deadlocks are caused by a number of situations, including:

- Transactions modify tables in different orders. There is a greater chance of deadlock between two transactions if one is processing in the sequence A–B–C while the other runs C–B–A.
- Transactions access tables using a nonclustered index. If the optimizer chooses a different nonclustered index for the same table for two different queries, a nonclustered index is not in the physical data sequence and the two processes are acquiring page locks in a random order. Therefore, there is a greater chance that one process will lock a page that the other needs.
- Transactions that use the keyword `holdlock` or use the `set isolation level` command to hold shared locks. When `holdlock` is appended to a select transaction it holds the shared lock for the remainder of the transaction. This increases the risk of deadlock.
- Transactions that require a long time to run. The longer a transaction runs, the more likely it is that another user will require a resource held by the transaction.

Action Restart the transaction that has been rolled back. To minimize future occurrences of deadlocks, use any of the following procedures that apply to your site.

Handling application errors

Each application should have deadlock handling routines. Refer to `dbsetuserdata` in the “Routines” chapter of the most recent version of the *Open Client DB-Library/C Reference Manual* (within the *Open Server 15.0, Open Client 15.0 and SDK 15.0* top-level topic on the Sybase Technical Library Product Manuals Web site) for a sample deadlock handling routine.

Using well-constructed transactions

Using transactions constructed to avoid deadlocks greatly reduces their occurrence. Some techniques for writing transactions that avoid deadlock include:

- Access tables in the same order in each transaction. Use coding conventions that require all transactions that access several tables to process them in the same order.
- Access tables via a clustered index when possible. If it is not possible to change a nonclustered index to a clustered index to minimize deadlocks, then trap the deadlock error in the application and provide appropriate recovery routines.
- If you are using `holdlock`, decide whether you really need to be using it. Use `holdlock` only when you require repeatable reads within a transaction.
- Avoid long-running transactions. Some ways to avoid long-running transactions are:
 - Never allow user interaction within a transaction.
 - Separate logical units of work into transactions. For example, acquiring a sequential key from a key table for use in an insert statement can be separated into transactions similar to the following (in the first set of commands, you acquire the key; in the second set of commands, you use the key to do the insert):

```
1> declare @key int
2> begin transaction
3> update key_table set key = key + 1
4> select @key = key
5> commit transaction
6> go
```

```
1> insert <mytable>
2> values (@key, ...)
3> go
```

	<ul style="list-style-type: none">• Check whether client applications, including third party tools, allow users to accidentally create long running transactions.
Additional information	Refer to “Lock management” in the most recent version of <i>Performance and Tuning: Monitoring and Analyzing</i> .
Version in which this error is raised	All versions

Error 1243

Severity	26
Message text	Couldn't find object to be released.
Explanation	<p>Adaptive Server uses logical locks for database pages and tables and address locks for a buffer containing OAM or allocation pages. When Adaptive Server no longer needs a lock, the server searches for it to release it.</p> <p>Error 1243 may occur because:</p> <ul style="list-style-type: none">• During the release of an address lock for a buffer containing OAM or allocation pages, Adaptive Server fails to find one of the buffers.• During the release of a logical lock for database pages and tables, the server fails to find the requested object on the lock chain.
Action	Shut down and restart Adaptive Server.

Note Since the unreleased lock may be an exclusive lock, there may be some objects you cannot access until you restart Adaptive Server.

If error 1243 occurs again, call Sybase Technical Support.

Version in which this error is raised	All versions
---------------------------------------	--------------

Error 1244

Severity	20
Message text	Could not find address lock at address 0x%lx while

trying to release it.

Explanation	Address locks are used by the database tasks to synchronize the access to Adaptive Server internal data structures. These tasks may get this error while relasing a lock that is already released or never acquired.
Action	This is a server internal error that results in rollback of the transaction and forces users to resubmit the transaction. Because this is an internal error, users cannot control this error.
Additional information	If this error persists, configuring a shared memory dump (CSMD) may be helpful in determining the root cause for this error. Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide. If necessary, contact Sybase Technical Support.
Version in which this error is raised	All versions

Error 1249

Severity	20
Message text	Process became runnable before lock it was waiting for was released.
Explanation	The Lock Manager maintains queues of lock requests to enable server tasks to obtain read, write, and update locks as needed. A task sleeps while waiting for a lock request to be granted. Error 1249 is raised when the task awakes but the Lock Manager detects an abnormal condition in the lock queue for that lock. Error 1249 is raised in the following states:

State	Meaning
1	The lock request was granted, but the lock is not at its expected position in the queue.
2	The lock received a signal to requeue, but it is not at its expected position in the queue.

Error 1249 is caused by an Adaptive Server problem and breaks the connection to the server. In a common scenario leading to this error, separate sessions open cursors on the same table, and the user aborts one session in the middle of an update; attempting the same transaction again raises the 1249 error.

Action	Call Sybase Technical Support.
--------	--------------------------------

Additional information	When you call Sybase Technical Support, have the information available that is listed in “Reporting errors” on page 4.
Version in which this error is raised	All versions

Error 1251

Severity	26
Message text	An in-use preallocated semaphore cursor was encountered.
Explanation	<p>A semaphore is the position at the head of a queue of locks, and is used by the server’s Lock Manager to ensure that tasks obtain valid and compatible locks. Error 1251 is raised when a task attempts to access a semaphore from a position in the queue that is already in use.</p> <p>1251 errors are due to a Adaptive Server locking synchronization problem and results in a stack trace.</p> <p>Some situations that can raise a 1251 error:</p> <ul style="list-style-type: none">• When you drop a column using alter table on an all-pages-locked (APL) table with a clustered index; and the server is configured for parallel processing;• If you create or alter a database in a device on AIX and the specified size is too large;• If a query in chained transaction mode invokes a Java method using iJDBC;• If there are stranded data extents in the log segment of a dedicated log.
Action	Call Sybase Technical Support for assistance; you may be able to upgrade to an Adaptive Server Enterprise version where the problem is resolved.
Additional information	<p>If the error was raised while creating or altering a database on AIX, note that the space actually available on a device is slightly smaller than the device size; specifying a slightly smaller size when you create or alter the database may help resolve the problem.</p> <p>If the error was raised while executing alter table drop column, disable parallel processing with sp_configure “max scan parallel degree” as a workaround.</p>

If the error was raised invoking a Java method, check your query. Commands like `begin tran`, `commit`, `rollback`, `set chained on/off` are not allowed in nested SQL. See *Java in Adaptive Server Enterprise* for details.

Before calling Sybase Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the query that raised the error.

Version in which this error is raised

All versions

Error 1265

Severity

20

Message text

Version 15.0.2 and later:

```
An illegal attempt was made to release an invalid lock
or a lock that is not owned by the process family
(lr=0x%p, lrmagic=%c).
```

Version 15.0.1 and earlier:

```
A lock manager routine expected a lock record as an
argument but was passed something else. lr=0x%x,
lrmagic=%c.
```

Explanation

This error occurs when the Adaptive Server Lock Manager is:

- Releasing a lock (state 1)
- Transferring a lock (state 2)
- Waiting for a lock (state 3)

and the lock record structure is not valid or the lock record belongs to another process.

This error is usually due to an Adaptive Server problem.

Action

Shutting down and restarting Adaptive Server usually clears the error. If this does not clear the error, or if the error persists, call Sybase Technical Support.

Additional information

Before calling Sybase Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the Output of `sp_lock` at the time the 1265 error occurred.

Version in which this error is raised All versions

Error 1279

Severity 16

Message text SQL Server has run out of locks on engine %d. Re-run your command when there are fewer active users, or contact a user with System Administrator (SA) role to reconfigure maximum engine freelocks or number of locks.

Explanation An Adaptive Server engine is a process running an Adaptive Server that communicates with other Adaptive Server processes via shared memory. An Adaptive Server running on a multiprocessor machine can have one or more engines.

Adaptive Server uses two configuration parameters to control lock use:

- number of locks controls lock use Server-wide
- max engine freelocks controls lock use for each Adaptive Server engine

Adaptive Server maintains a list of Server-wide free locks and a list of per-engine free locks. When locks are needed for an engine, Adaptive Server moves a block of free locks from the Server-wide free lock list to the free lock cache for that engine.

Error 1279 occurs when Adaptive Server runs out of locks on an engine and is unable to move any more locks from the Server-wide free list but there is another engine with some free locks. (Error 1204 occurs when Adaptive Server completely runs out of locks.) The error message includes the number of the engine that has run out of locks.

Action Run your command again later when there are fewer active users or increase the number of locks available to the engine. If you want to increase the resources, you can do one of the following:

- Increase the total number of locks available Server-wide by using `sp_configure` to increase the value of the number of locks configuration parameter.

- Decrease the maximum number of locks available for each engine by using `sp_configure` to decrease the value of the `max engine freelocks` configuration parameter. By decreasing the value of this parameter, you make more locks available to the engine for which you are getting the error.

Note Increasing the value of the `number of locks` parameter increases the memory resources used by Adaptive Server. Make sure your Adaptive Server memory is configured appropriately if you change this value.

Additional information

Refer to “Lock management” in the most recent version of *Performance and Tuning: Monitoring and Analyzing* for information about the Adaptive Server Lock Manager.

Refer to the *Reference Manual: Procedures* for information about `sp_configure` and configuration parameters.

Refer to “Memory management: in the most recent version of the *Performance and Tuning: Monitoring and Analyzing* for information about Adaptive Server memory management.

Version in which this error is raised

All versions earlier than 12.5

Sort Manager Errors (1500s)

This section contains error messages for the Adaptive Server Enterprise Sort Manager, which determines whether a given sort should be performed in parallel or in serial.

Error	Page
Error 1501	215
Error 1505	217
Error 1506	218
Error 1508	219
Error 1509	220
Error 1510	221
Error 1514	223
Error 1530	224
Error 1531	224

Error 1501

Severity	20
Message text	Sort failure
Explanation	

This error occurs during index creation when Adaptive Server fails to complete a sort operation. There are a number of different states for this error, each indicating a different reason for the failure:

State	Meaning
2	When creating a high level index structure, Adaptive Server is unable to build a parent node in the ancestor list.
3	When creating a high level index structure, Adaptive Server is unable to build a parent node from the root of the tree.
4	When creating a high level index structure, Adaptive Server is unable to build a parent node from the new root of the ancestor tree.

State	Meaning
6	The server is adding the results of an internal sort of a special leaflet node to an external merge in progress, and a single internal merge step has a merge order greater than 8.
7	The server is adding the results of an internal leaf merge node and the previous merge step caused the merge order to be greater than 8.
8	Adaptive Server is starting a non-sort of already sorted data to create a clustered index, and it cannot find a page to an index that was supposed to have been built for the sorted data.
9	The server is unable to allocate and initialize space for managing the sort's buffers.
10	The server is building indexes and the length of a row in an overflow page is either zero or greater than the size of a page.

When this error occurs, no index is created.

The configuration parameter `number of sort buffers` controls how Adaptive Server uses memory during sorting. `number of sort buffers` controls the sort buffer size; its default value is 500. Changing the value of `number of sort buffers` can cause error 1501.

Error 1501 can also be caused by an Adaptive Server problem.

Action

To resolve the error:

- 1 Determine the current value of `number of sort buffers`:


```
1> sp_configure "number of sort buffers"
2> go
```
- 2 Change the value, perhaps back to the default value of 500 (recommended except when creating indexes in parallel):


```
1> sp_configure "number of sort buffers",
<new_value>
2> go
```

Additional information

Refer to “create index” in *Reference Manual: Commands* for information about creating indexes.

Refer to “Setting Configuration Parameters” in the *System Administration Guide: Volume 1* for information about configuration parameters.

Version in which this error is raised

All versions

Error 1505

Severity	14
Message text	Create unique index aborted on duplicate key. Primary key is '%S_KEY'.
Explanation	<p>This error occurs when you try to create a unique index on a column or columns containing at least two rows with the same index value. Each time a unique index is created, Adaptive Server checks for duplicate index values (if data already exists), and displays this message if it finds any.</p> <p>When this error occurs, no index is created.</p> <p>A command of the following type generates a 1505 error if there are duplicate values on <code>col1</code>, <code>col2</code>, and <code>col3</code>:</p> <pre>create unique index <index_name> on <table_name>(<col1>, <col2>, <col3>)</pre>
Action	<p>Decide whether you want to allow or prevent duplicate keys in the table where the index will be created. Use one of the following procedures:</p> <ul style="list-style-type: none"> • If you decide to allow duplicate keys, do not use the keyword <code>unique</code> when you create the index. • If you decide to prevent duplicate keys, delete rows that have duplicate index values. <p>If you choose to delete duplicate keys, use the following procedure. For a brief period of time some rows from the table exist only in a temporary table. If a power failure or other interruption occurs, you could lose the temporary table and the rows in it. Therefore, before you begin this procedure, make a backup copy of the table or modify the procedure to use a table you create for this purpose instead of a temporary table.</p> <ol style="list-style-type: none"> 1 Find the rows that caused the error: <pre>1> select * from <table_name> 2> group by <col1>, <col2>, <col3> having count(*) > 1 3> go</pre> <p>The query result contains all the rows causing the 1505 error.</p> 2 Put one instance of each distinct row from step 1 into a temporary table: <pre>1> select distinct * into #<temp_table_name> 2> from <table_name> 3> group by <col1>, <col2>, <col3> having count(*) > 1 4> go</pre>

- 3 Some of the rows may contain identical values in the column or columns to be indexed and different values in other columns. Use the query below to find these rows:

```
1> select * from #<temp_table_name>
2> group by <col1>, <col2>, <col3> having count(*) > 1
3> go
```

Examine the individual rows to decide which rows should be deleted. Delete these rows from the temporary table now. The rows you delete in this step will be deleted from the original table in step 4.

- 4 Delete all the rows which caused the 1505 error from the original table:

```
1> delete <table_name> from #<temp_table_name>
2> where <table_name.col1> = #<temp_table_name.col1>
3> and <table_name.col2> = #<temp_table_name.col2>
4> and <table_name.col3> = #<temp_table_name.col3>
5> go
```

- 5 Each row in the temporary table should have unique values for the columns in question. Add these rows back into the original table:

```
1> insert into <table_name>
2> select * from #<temp_table_name>
3> go
```

- 6 Now you can create the unique index on the table.

Additional information

Refer to “create index” in the *Reference Manual: Commands* for information about indexes.

Version in which this error is raised

All versions

Error 1506

Severity

20

Message text

```
Sort in database '%.*s' encountered an internal error
in logging state. Error code : %d. Transaction ID: (%s).
Object ID of target table: (%ld)
```

Explanation

It is an internal error during a sort. Adaptive Server encountered an issue trying to write log records for the data in syspartitions rows.

Action	Resubmit the query. If this error is raised more than once, please contact Sybase Technical Support.
Additional information	Before you call Sybase Technical Support, set up a configurable shared memory dump (CSMD), which may be helpful in determining the root cause for this error. Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 1508

Severity	14
Message text	Create index aborted on duplicate rows. Primary key is 'S_KEY'.
Explanation	This error occurs when you try to create a clustered index on a table that contains duplicate rows.
Action	Decide whether you want to allow or prevent duplicate rows in the table where the clustered index will be created. Use one of the following procedures.

Allowing duplicate rows

Create the index with the `allow_dup_row` option:

```
1> create clustered index <index_name> on <table_name>
2> with allow_dup_row
3> go
```

Subsequent update and insert commands can create duplicate rows after the `allow_dup_row` option is set.

Preventing duplicate rows

Create the index with the `ignore_dup_row` option:

```
1> create clustered index <index_name> on <table_name>
2> with ignore_dup_row
3> go
```

Existing duplicate rows will be deleted from the table as the index is created. Subsequent attempts to enter a duplicate row with `insert` or `update` are ignored and the `insert` or `update` is cancelled with an informational message. If the `insert` or `update` is part of a larger transaction, the rest of the transaction will continue as usual.

Note `ignore_dup_row` and `allow_dup_row` are mutually exclusive. Using them together in the same `create index` statement causes error 1916 .

Version in which this error is raised

All versions

Error 1509

Severity 20

Message text Row compare failure

Explanation This error occurs when Adaptive Server compares two rows during a sort and the compare fails because a row was being compared against itself.

Error 1509 can occur during normal processing or when you try to create an index on a table and the operation fails because of data corruption or an Adaptive Server problem.

This is a serious error as it indicates possible corruption.

Action

- 1 Run `dbcc checktable` and `dbcc tablealloc` on the table being accessed at the time of the error to check for corruption.
- 2 If there is no corruption, call Sybase Technical Support.
- 3 If there is corruption, refer to the write-ups in this manual for other error messages that appear. Run `dbcc checktable` and `dbcc tablealloc` again to determine whether the corruption has been cleared.

- 4 If corruption still exists after eliminating all error messages, copy your data out and back in using the procedure in “Rescuing data from a corrupted table” in *Troubleshooting and Disaster Recovery*, or restore your database from clean backups using the procedure below.

Warning! Some data might be lost on this page if you recover your table using `bcpl` or `select into` (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

If you choose to restore your database from clean backups, do the following:

- 1 Drop the database. If the drop fails, follow the instructions in “Dropping a database when drop database fails” in *Troubleshooting and Disaster Recovery*.
- 2 Create a database for load. Make sure the database you create has sizes as least as large as those in `sysusages` for the original database (and that all other `sysusages` values match the original values). For more information about creating a database to load, refer to “create database” in *Reference Manual: Commands*.
- 3 Load the database from backup. Refer to “load database” in *Reference Manual: Commands*.
- 4 Run the online database command to make the database available for use.

Additional information

If you need to call Sybase Technical Support, have the information ready that is listed in “Reporting errors” on page 4, including the output from `dbcc checktable`, `dbcc tablealloc`, and `sp_configure`.

Version in which this error is raised

All versions

Error 1510

Severity

17

Message text

Sort failed: Out of space in database '%.*s'

Explanation Adaptive Server uses a clustered index to sort rows so that their physical order is the same as their logical (indexed) order. When creating a clustered index, Adaptive Server needs a considerable amount of space, 120 to 150 percent of the size of the table, because it needs to sort the data in the table first and then create the clustered index on the sorted data. This space should be available in the database or on the segment where you want to create your index.

Error 1510 occurs when you try to create an index and one of the following is true:

- There is not enough space in the database to create the index. Usually this error occurs with error 1105.
- No more locks are available.

Unfortunately, there is no straightforward way to identify which problem caused the error.

Action Select the appropriate action from the sections below.

Not enough space in the database

Use either of the following two strategies to correct this error:

- Use the instructions from [Error 1505](#) to increase the space available in your database.
- Move the table where the index is to be created to a different, larger, user-defined segment, and then create the index. Refer to “Creating and Using Segments” in the *System Administration Guide: Volume 2* for instructions.

- a Determine whether you have enough space available on the new segment by typing:

```
1> sp_helpsegment <segment_name>
2> go
```

- b Verify that the database where you want to create your index has enough space allocated on the device or devices that contain the segment from the above query:

```
1> sp_helpdb <database_name>
2> go
```

No more locks available

If the 1510 error occurs because Adaptive Server has no more locks available, use `sp_configure` to increase the number of locks:

```
1> sp_configure "number of locks", <new_value>
2> go
```

Shut down and restart Adaptive Server to initialize the new configuration parameter value.

Refer to “Setting Configuration Parameters” in the *System Administration Guide: Volume 1*, and “Configuring Memory” in the *System Administration Guide: Volume 2* for more information.

Additional information

Refer to “create index” in the *Reference Manual: Commands*.

Version in which this error is raised

All versions

Error 1514

Severity

21

Message text

Page allocated to sort found to be busy. Page number %ld. %S_BUF.

Explanation

During index creation, Adaptive Server allocates and uncaches pages to sort the table’s rows. The sort routine double checks whether a page allocated for the sort is in the data cache. If it is in the cache, the page cannot be used, the sort operation fails and error 1514 is raised.

Failure to sort means that the index creation fails; consequently 4313 errors often accompany error 1514.

Action

If error 1514 is raised in an isolated case during index creation on a specific table, run `dbcc checktable` on the table. If errors are reported, try to fix them using the sections in this manual that describe those errors, then use one of these options to recover from error 1514:

- Retry index creation.
- Shut down and restart Adaptive Server, and retry index creation.
- Select the table data into a new table, drop the old table, and rename the new table to the old table name. Create the desired index on the table.
- Bulk copy the affected table out, drop and re-create the table, and bulk copy back in. Create the desired index on the table.

If the problem persists, or is not isolated to one table, contact Sybase Technical Support for assistance.

Version in which this error is raised

All versions

Error 1530

Severity	16
Message text	Create index with sorted_data was aborted because of row out of order. Primary key of first out of order row is '%S_KEY'
Explanation	<p>The sorted_data option speeds creation of an index when the data in the table is already in sorted order. The speed increase becomes significant on tables larger than 1GB.</p> <p>Error 1530 occurs when you try to create an index with the sorted_data option on a column or columns containing data that is not in sorted order.</p> <p>When this error occurs, no index is created.</p> <hr/> <p>Note The sorted_data option speeds index creation only for clustered indexes or unique nonclustered indexes. Creating a nonunique, nonclustered index with the sorted_data option may succeed, but there is no improvement in performance.</p> <hr/>
Action	<p>To create an index after getting a 1530 error, do one of the following:</p> <ul style="list-style-type: none">• Sort the data (using your standard operating system procedures), bcp it into the table, and then create the index with the sorted_data option.• Create the index without the sorted_data option.
Additional information	Refer to “create index” in the <i>Reference Manual: Commands</i> for information about indexes.
Version in which this error is raised	All versions

Error 1531

Severity	16
Message text	The sorted_data option cannot be used for a nonclustered index if the keys are not unique within the table. Create index was aborted because of duplicate keys. Primary key is '%S_KEY'.

Explanation This error occurs when you try to create a nonunique, nonclustered index with the `sorted_data` option on a column or columns containing rows with the same index value. If this situation exists, a statement of the following type will fail:

```
1> create index <index_name>
2> on <table_name>(<col1>, <col2>, <col3>)
3> with_sorted_data
4> go
```

When error 1531 occurs, no index is created.

Action To avoid the 1531 error, do one of the following:

- Do not use the `sorted_data` option, or
- Delete rows with duplicate index values.

Deleting rows with duplicate index values

If you decide to delete rows with duplicate index values, use the following procedure.

Note For a brief time, some rows from the table will exist only in a temporary table, and if a power failure or other interruption occurs, you could lose the temporary table and the rows in it. To protect yourself against data loss, you can make a backup copy of the table before you begin the procedure. Alternatively, you can modify the procedure to use a specific table instead of a temporary table.

1 Find the rows that caused the 1531 error:

```
1> select * from <table_name>
2> group by <col1>, <col2>, <col3>
3> having count (*) > 1
4> go
```

2 Put one instance of each distinct row from step 1 into a temporary table:

```
1> select distinct * into #<temp_table_name>
2> from <table_name>
3> group by <col1>, <col2>, <col3> having count (*) > 1
4> go
```

3 Some of the rows may contain identical values in the column or columns to be indexed and different values in other columns. Use the following query to find these rows:

```
1> select * from #<temp_table_name>
2> group by <col1>, <col2>, <col3> having count (*) > 1
```

```
3> go
```

- 4 Examine the individual rows to decide which rows should be deleted. Delete these rows from the temporary table now. The rows you delete in this step will be deleted from the original table in step 5.

- 5 Delete all the rows which caused the 1531 error from the original table:

```
1> delete <table_name> from #<temp_table_name>
2> where <table_name.col1> = #<temp_table_name.col1>
3> and <table_name.col2> = #<temp_table_name.col2>
4> and <table_name.col3> = #<temp_table_name.col3>
5> go
```

- 6 Now each row in the temporary table should have unique values for the columns in question. Add these rows back into the original table:

```
1> insert into <table_name>
2> select * from #<temp_table_name>
3> go
```

- 7 Create the index with the `sorted_data` option.

Additional information

Refer to “create index” in the *Reference Manual: Commands* for information about indexes.

Version in which this error is raised

All versions

Initialization Errors (1600s)

Each database device or file must be prepared and made known to Adaptive Server before it can be used for database storage. This process is called initialization. This section contains error messages for Adaptive Server Enterprise initialization routines.

Error	Page
Error 1601	227
Error 1602	229
Error 1603	230
Error 1605	231
Error 1608	234
Error 1613	235
Error 1621	236
Error 1622	238
Error 1623	238

Error 1601

Severity	21
Message text	There are not enough 'user connections' available to start a new process. Retry when there are fewer active users, or ask your System Administrator to reconfigure SQL Server with more user connections.
Explanation	<p>Error 1601 occurs when all user connections are in use and a request is made to create another Adaptive Server process.</p> <p>Error 1601 occurs when all user connections are in use and a request is made to create another Adaptive Server process.</p> <p>Error 1601 is typically caused when Adaptive Server is not configured with enough user connections to support all requests.</p>
Action	Perform the actions from the appropriate sections below.

Too few user connections (all operating systems)

Use `sp_configure` to increase the value of the number of user connections configuration parameter:

- 1 Determine the current number of user connections configured:

```
1> sp_configure "number of user connections"  
2> go
```

- 2 Set the new value for user connections:

```
1> sp_configure "number of user connections",  
2> <new_value>  
3> go
```

- 3 There is no need to restart Adaptive Server in version 15.0 and later.

If you are running a version of Adaptive Server Enterprise earlier than 15.0, restart Adaptive Server to activate the changes.

Depending on the exact configuration of your Adaptive Server, each user connection typically requires between 40K and 60K of memory, whether or not it is in use. Gradually increase the number of user connections until these errors no longer occur. When you change the value of the number of user connections configuration parameter, you may need to increase the total memory configuration parameter. If you have increased the stack size or default network packet size configuration parameters, each user connection will require more memory. Refer to the Adaptive Server Enterprise performance and tuning guides and the chapter “Configuring Memory” in the *System Administration Guide: Volume 2* for information about how Adaptive Server uses memory.

For more information on determining the number of user connections your Adaptive Server needs, refer to “number of user connections” in the chapter “Setting Configuration Parameters” in the *System Administration Guide: Volume 1*.

Note This error often appears in the Adaptive Server error log without the 1601 error number: “kernel: no pss structure available for new process.” Take the above action whether or not the error number is displayed.

Additional information

You can obtain more details about user connections by executing the command `sp_monitorconfig 'user connections'`.

Version in which this error is raised

All versions

Error 1602

Severity	21
Message text	Unable to initialize network %d

Note This error may be caused by a hardware problem.

Explanation This error occurs during start-up when Adaptive Server is unable to initialize a network connection specified in your network addressing (*interfaces* or *sql.ini*) file.

Note Throughout this write-up, your network addressing file is referred to as your *interfaces* file.

Error 1602 can occur as a result of any of the following problems in your *interfaces* file:

- Duplicate entries
- Incorrect network address
- Other incorrect information such as incorrect host name or illegal port number
- Windows:
 - The WINSOC libraries or the protocol on which you want to run (that is, IPX/SPX) are not loaded.
 - You try to bring up the Adaptive Server on a bad default server name (Windows uses default values).

Error 1602 can also occur as a result of changes in your network that make entries in your *interfaces* file unavailable to Adaptive Server.

Action Check your *interfaces* file to make sure everything is correct. If any information is incorrect or duplicate entries exist, use one of the following to make corrections to your *interfaces* file:

- dsedit
- dscp

If you do not find any incorrect information in your *interfaces* file, ask your network administrator to verify (at the operating system level) the entries in the *interfaces* file to make sure all network connections are working properly.

Additional information For information about modifying the interfaces file for your platform, refer to the Adaptive Server installation and configuration guide.

If you need to call Sybase Technical Support, have your error log output and a copy of your *interfaces* or *sql.ini* file ready before you call.

Version in which this error is raised All versions

Error 1603

Severity 21

Message text Process priority %d invalid or no process slots available

Explanation When starting a new process, Adaptive Server checks that:

- The process is being started at the correct priority. Valid priorities range from 0 (highest) to 7 (lowest).
- There are enough connections (slots) available.

Error 1603 occurs when the new process is requested at an invalid priority, or when not enough connections are available to start the process.

Error 1603 occurs with the following states:

State	Meaning
1	Unable to create a new, internal server process.
2	Unable to create a process to handle a new user connection (for example, an isql connection) to the server.

Action Check the following parameters, and increase the values as needed:

```
number of user connections
number of remote connections
```

Also check any applications currently accessing the server. The number of connections needed by an application depends entirely on how the application is programmed. You may need to shut down and correct any applications that are using excessive connections.

Process priority

Check the value of process priority reported in the message text. If the value does not fall in the valid range between 0 and 7, check the priority at which applications are accessing the server.

Additional information

number of user connections and number of remote connections are static parameters. You must restart the server after you change their values.

Refer to the chapter “Setting Configuration Parameters” in the *System Administration Guide: Volume 1* for details about parameters that control new connections.

You can also obtain more details about user connections by executing the command `sp_monitorconfig 'user connections'`.

Version in which this error is raised

All versions

Error 1605

Severity

21

Message text

Failed to open virtual socket for new connections

Note This error may be caused by a hardware problem.

Explanation

The Adaptive Server Enterprise Network Handler listens on the ports specified in your network addressing (*interfaces* or *sql.ini*) file. When a connection comes in, the Network Handler opens another socket to redirect that client, clearing the listener port for the next connection, hence the term “virtual” socket. The connection appears, from the outside, to be connected to the *interfaces* port, when it is actually connected at a different address to which there is a pointer.

Error 1605 occurs when a failure occurs in the process of opening the new socket. There are a number of things that can cause this:

- Exceeding the number of configured user connections
- Exceeding the number of configured sockets
- Network configuration problems such as using the same IP address for two different clients

Error 1605 error can be raised at:

- Startup time, when the network is initialized.
- Creation of a process, when a new connection request comes in and Adaptive Server is unable to handle it.
- The forced termination of a process, when a new connection request comes in and Adaptive Server is unable to handle it.

The information below includes several internal configuration parameters. You will not be changing these directly, but they are included to help explain how to calculate values for `sp_configure` parameters that relate to connection handling. Following are the kernel and configuration parameters used to calculate `cnssocket` (the maximum number of virtual sockets):

- `kimaxfd` – the maximum number of configured per process file descriptors that a process can use at a given time. The error log message “Using ## file descriptors” shows what this value is at boot time.
- number of remote sites – this configuration parameter specifies the maximum number of remote sites.
- number of user connections – this Adaptive Server configuration parameter specifies the maximum number of user connections allowed.
- `cnssocket` – the maximum number of virtual sockets, calculated as:

```
cmaxnetworks (maximum # of user connections configured)
+ cusrconnections (the number of user connections configured)
+ cfgsitebuf (the number of remote sites configured)
+ XTRA_SOCKETS (defined as 2 in 11.0.x)
```

During start-up, if the value for `cnssocket` is greater than the value for `kimaxfd`, then the value of `cnssocket` will be set to the value of `kimaxfd`, and the message:

```
kistartup: number of virtual sockets set to ###
```

is printed to the log where `###` is the value of `kimaxfd/cnssocket`.

Adaptive Server recalculates the value for `cnssocket` when you use `sp_configure` to change related parameter values.

Action

Check the error that precedes the 1605 error. If that error is one of the following:

```
nopen: No virtual sockets available
nopen: no room in socket table
```

then the value for the Server configuration block parameter *cnssocket* has been exceeded. To solve the problem in this case, increase user or network connections and possibly increase the maximum number of file descriptors per process:

- 1 Determine the current value for *cnssocket*.

Use `dbcc resource` (this command requires `sybase_ts_role` and `dbcc tracecon(3604)` to determine the current configuration for sockets.

- 2 Determine the current value for the maximum number of allowed network connections:

```
1> select @@max_connections
2> go
```

`@@max_connections` is the maximum number of file descriptors that the operating system allows for your server's engine processes, minus a fixed number of descriptors needed by Adaptive Server. Refer to the chapter "Managing Multiprocessor Servers" in the *System Administration Guide: Volume 2* for details about the `@@max_connections` calculation.

- 3 Determine the current value for number of devices (the number of database devices that Adaptive Server can use, excluding dump devices):

```
1> sp_configure "number of devices"
2> go
```

- 4 Determine an appropriate value for `@@max_connections`, based on the number of users expected to be logged on at any given time, and the number of devices. Remember that each device takes up a user connection.
- 5 Compare the appropriate value for sockets to the current value for sockets which you determined in step 1. (Remember that `maximum number of virtual sockets = max number network listeners + number of user connections + number of remote sites + 2`.)

If the current value is too low, use `sp_configure` to increase the value of number of user connections and/or number of remote sites so that the calculated value for sockets will be correct. If the error persists, call Sybase Technical Support.

Note Under UNIX, you may need to raise the maximum number of file descriptors at the operating system level for the Adaptive Server if the new value for `cnsocket` exceeds the maximum number of per process file descriptors times the number of online engines for 11.0.x and later versions.

Refer to “number of user connections” in the chapter “Setting Configuration Parameters” in the *System Administration Guide: Volume 1* for information about managing the maximum number of user connections available on an Windows server.

Additional information	Refer to “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for information about configuration parameters. Refer to the <i>Reference Manual: Procedures</i> for information about <code>sp_configure</code> . For parameter information specific to your platform, refer to the Adaptive Server installation and configuration guide for your platform. If you need to call Technical Support, have the information listed in “Reporting errors” on page 4 available, including the output of <code>dbcc resource</code> and <code>sp_configure</code> .
Version in which this error is raised	All versions

Error 1608

Severity	18
Message text	A client process exited abnormally, or a network error was encountered. Unless other errors occurred, continue processing normally.
Explanation	This error occurs when a client process stops without informing Adaptive Server and a subsequent attempt by Adaptive Server to send results to this client fails.

Error 1608 is usually not serious and may be viewed as notification that a client process no longer exists. Some of the possible reasons a client process can disappear are:

- The client application was killed or ended abnormally
- The client machine was powered off or was restarted
- There is excessive network traffic
- The network has been shut down

Note Using Ctrl-C (on operating systems other than Windows) to break out of an application should not cause this error.

Action Do not be concerned if this error occurs only sporadically. However, if the error occurs frequently and continuously, or frequently for short periods of time, it may be a result of excessive network traffic or network problems.

If you suspect network problems, have your network administrator run diagnostics on the network between the client and the Adaptive Server machines to isolate the problem.

Additional information Version 12.5.3 and later include an enhancement to error 1608, which now displays the host name and login name when a client connection is disconnected abnormally from Adaptive Server. The extended error information provides additional diagnostics. The System Administrator can now analyze and confirm the cause of the disconnect. The error message appears as follows:

```
00:00000:00011:2007/09/21 16:52:36.60 server Error:1608, Severity: 18, State:4
```

```
00:00000:00011:2007/09/21 16:52:36.60 server A client process exited
abnormally, or a network error was encountered. Unless other errors occurred,
continue processing normally.
```

```
00:00000:00011:2007/09/21 16:52:36.60 kernel extended error information:
hostname: TELMOW2KADS login: sa
```

Version in which this error is raised All versions

Error 1613

Severity

21

Message text Could not close network %d connection for server process %d.

Explanation This error occurs when Adaptive Server is unable to close a network connection for a Server process.

Error 1613 can be caused by:

- System parameters being set too low
- Network problems
- Problems specific to an operating system (see the next section “Action” for a Digital Unix Tru64 version 5 example)

This is not a serious problem unless it occurs frequently.

Note This error may be caused by a hardware problem.

Action If you are running Adaptive Server on Digital Unix Tru64 version 5, you will need the appropriate operating system patch that corrects this problem. Go to www.sybase.com/support and search for TechNote 1012739 for the list of patches and patch kits available at the time of publication. As always, check with the operating system vendor for the latest patches.

On other platforms, check the values for your system parameters that affect network use:

- UNIX – these might include the limit on the maximum number of file descriptors (limits the number of sockets that can be open at one time) and the per-process limit on file descriptors.
- Windows – these might include NumForwardPackets and ForwardBufferMemory.

If you suspect network problems, have your network administrator run diagnostics on the network between the client and the Adaptive Server machines to isolate the problem.

Version in which this error is raised

All versions

Error 1621

Severity

18

Message text	Type '%c' not allowed before login.
Explanation	<p>Adaptive Server uses Tabular Data Stream™ (TDS) packets to communicate with clients. A token is attached to every TDS packet to differentiate between different types of packets. During login, Adaptive Server expects a client to supply either the TDS_LOGIN or TDS_ECHO token.</p> <p>Error 1621 occurs when a client sends a token other than TDS_LOGIN or TDS_ECHO at login. Possible causes include:</p> <ul style="list-style-type: none"> • Non-printing characters in the login script. • Another process trying to access the port used by the server. • An incompatible database driver; for example, an unsupported ODBC driver.
Action	<p>To see if other processes are interfering with the port, use <code>netstat</code> (UNIX) or Task Manager (Windows) and check port status. You can also start the server using a different port number to remove any possible interference.</p> <p>Check client applications to make sure logins are using a supported database driver. On Windows, you can use the Drivers tab under ODBC Data Source Administrator to see the available drivers.</p> <p>Check the TDS login packets being sent to Adaptive Server. To identify the affected logins, you can monitor packets by using:</p> <ul style="list-style-type: none"> • Trace flags 3605 and 4013 (3605 sends the output to the error log, and 4013 generates the login record) <hr/> <p>Warning! Use of trace flag 4013 may generate a large number of messages in the error log.</p> <hr/> <ul style="list-style-type: none"> • The <code>ribo</code> utility (version 12.0 and later). <p>Correct your login scripts to remove non-printing characters or control characters, and retry the connection.</p>
Additional information	Refer to “Starting Adaptive Server with trace flags” in the chapter “Other Useful Tasks” in the most recent version of <i>Troubleshooting an Disaster Recovery</i> for details on using trace flags.
Version in which this error is raised	All versions

Error 1622

Severity	18
Message text	Type '%c' not implemented.
Explanation	Adaptive Server uses TDS (Tabular Data Stream) packets during communication with clients to differentiate between different types of packets. A token is attached to every TDS packet. Error 1622 occurs when a client sends an invalid TDS token.

Valid packet header tokens include:

- TDS_LANGUAGE
- TDS_RPC
- TDS_ECHO
- TDS_LOGOUT

In the message output, “%c” is the token name causing the 1622 error.

Error 1622 occurs with the following states:

State	Meaning
1	Error 1622 occurs with state 1 if the packet header cannot be converted to one of the tokens listed above.
2	Error 1622 occurs with state 2 if the header token is correct but the data token is bad.

Action	Check your network to make sure there are no unexpected characters being sent to Adaptive Server because of noise on the line or hardware problems. If you have confirmed that unexpected characters are not occurring, call Sybase Technical Support.
--------	--

Version in which this error is raised	All versions
---------------------------------------	--------------

Error 1623

Severity	21
Message text	Warning: dataserver -p option fails. Please check for state '%d' to see where it fails.

Explanation This error occurs when Adaptive Server is unable to create a new password for a user with “sso” role in response to the `dataserver` command with the `-p` option.

This error is raised with three possible states:

State	Meaning
1	Login passed with the <code>-p</code> option does not exist in <code>syslogins</code> .
2	Login exists but does not have the “sso” (system security officer) role.
3	Call to the routine that sets the passwords failed. This routine generates its own errors to help diagnose the problem.

Action The actions for recovering from error 1623 are listed below, by state.

States 1 and 2

If you did not remove the “sso” role from the “sa” login, use “sa” in the `dataserver` command in your `runserver` file. If you did remove the “sso” role from the “sa” login, choose the option that is appropriate from the following choices.

- If you can, log in to Adaptive Server as any user.

Determine which login names have “sso” role:

```
1> use master
2> go
```

```
1> select l.name, s.name from sysloginroles slr,
2> syslogins l,sysssrvroles s
3> where slr.srid = 1 and slr.suid = l.suid and
4> slr.srid = s.srid
5> go
```

- If you cannot log into Adaptive Server as any user, and you have a historical dump of the `master` database that includes a user with the “sso” role and for which you know the password, create a new `master` device, load that dump and determine whether you can log into it.

Refer to “master database is corrupt and Adaptive Server starts,” “Valid dump of the `master` database exists” in the chapter “Adaptive Server System Recovery” in the most recent version of *Troubleshooting and Disaster Recovery* for instructions on how to re-create your `master` database.

- If you cannot log into Adaptive Server as any user and do not have a historical dump of the master database, refer to “master database is corrupt and Adaptive Server does not start,” “Valid dump of the master database does not exist” in the chapter “Adaptive Server System Recovery” in the most recent version of the *Troubleshooting and Disaster Recovery* for instructions on how to re-create your master database.

State 3

Other error messages will be listed. Follow the directions in this manual for recovering from those errors. If no directions exist, call Sybase Technical Support.

Version in which this error is raised

All versions

Create Utilities Errors (1700 - 1900s, 2100s, 2700s, 12800s)

This section contains error messages for Adaptive Server create operations (create table, create database, and so on).

Error	Page
Error 1702	242
Error 1708	242
Error 1732	243
Error 1803	244
Error 1808	245
Error 1809	247
Error 1810	248
Error 1813	248
Error 1820	249
Error 1902	250
Error 1903	251
Error 1916	252
Error 1928	253
Error 2110	254
Error 2714	255
Error 2729	256
Error 2753	257
Error 2762	258
Error 12818	260
Error 12842	261
Error 12881	262

Error 1702

Severity	16
Message text	<code>%s failed because definition of column '%.*s' in table '%.*s' exceeds the limit of %d columns per table.</code>
Explanation	<p>This error occurs when you try to create a table with more than the allowed number of columns. The maximum number of definable columns per table is 250 in version 12.0.x and earlier, and 1024 in version 12.5 and later (254 for variable width all-pages-locked (APL) tables).</p> <hr/> <p>Note Column names must follow the rules for identifiers. They must be unique within a given table, but you can use the same column name in different tables in the same database.</p> <hr/>
Action	<p>Reduce the number of columns in the table you are trying to create to or less than the limit. Refer to “create table” in the <i>Reference Manual: Commands</i> for details.</p>
Version in which this error is raised	All versions

Error 1708

Severity	16
Message text	<code>Warning: Row size (%d bytes) could exceed row size limit, which is %d bytes.</code>
Explanation	<p>This is an information message printed when creating a table or adding and modifying table columns using the <code>alter table</code> command). This message is printed when the size of the data row of the table being created or altered exceeds the maximum data row size allowed in the server (maximum row size depends on the lock scheme of the table).</p> <p>This error does not terminate the execution of the statement and allows the table to be created or altered. But, when an application tries to insert or update a row that is larger than the displayed size limit, that operation will fail with an error such as 511.</p>
Action	<p>Carefully design the schema so that the data row size is under the maximum data row size, which depends on lock scheme, allowed by the server.</p>

Version in which this error is raised Version 12.5 and later

Error 1732

Severity 20

Message text Failed to reserve/initialize allocation hints space in table '%.*s'.

Explanation Each table, index, and text chain has one or more object allocation map (OAM) pages allocated to the table or index. These OAM pages store pointers to each allocation unit that contains pages for the object. The object's first OAM page stores allocation hints, indicating which OAM page has information about allocation units with free space.

Error 1732 is raised when Adaptive Server is allocating pages for a new table, but is unable to create allocation hints space for the table.

Error 1732 occurs with the following states:

State	Meaning
1	Unable to create allocation hints space for the table.
2	Unable to initialize allocation hints space for the table.
3	Unable to create allocation hints space for a text page chain.
4	Unable to initialize allocation hints space for a text page chain.

The table named in the message may be a user table, or a temporary table the server creates when executing a stored procedure or other query. Error 1732 can be due to an Adaptive Server problem.

Action This is not a serious error. Reconnect to the server and retry the query which raised the error; it should run normally.

Contact Sybase Technical Support if the problem persists.

Additional information Have the information listed in “Reporting errors” on page 4 before calling Sybase Technical Support.

Version in which this error is raised All versions

Error 1803

Severity 17

Message text CREATE DATABASE failed. Could not allocate enough disk space for a new database on the disks named in the command. Total space allocated must be at least %d Mbytes (%ld 2048-byte pages) to accommodate copy of Model Database.

Explanation When you issue a create database command, Adaptive Server:

- Assigns space to the database on the specified database devices. If you use the default keyword or if you omit the `on` clause altogether, Adaptive Server puts the database on one or more of the default database devices specified in `master..sysdevices`.
- Makes a copy of the `model` database in the new database space, creating the new database's system tables. The new database thus inherits all the changes you have made to the `model` database, including the addition of user names and objects.

Error 1803 is raised when Adaptive Server is unable to allocate the space required for the `model` database on the specified devices.

Action Make sure that the available space on the database device you specified (or the available space on the default device) will accommodate the `model` database.

To check the size of `model`, use the following command:

```
1> sp_helpdb model
2> go
```

To check the space available on the device, first use the following commands to check the space already allocated to other databases on the device:

```
1> use master
2> go
```

Then enter the following commands, depending on the version of Adaptive Server Enterprise you are using:

- Versions earlier than 15.0:

```
1> select dbid, size, phname "physical device"
2> from sysusages, sysdevices
3> where name = '<device_name>'
4> and vstart between low and high
5> compute sum(size)
6> go
```


For example:

```
1> select dbid, size, phyname "physical device"
2> from sysusages, sysdevices
3> where name = 'sd5f'
4> and vstart between low and high
5> compute sum(size)
6> go
```

```
dbid    size    physical device
-----  -
15      17920  /dev/rsd5f
16      20480  /dev/rsd5f
17       7680  /dev/rsd5f
18      20480  /dev/rsd5f
21       5120  /dev/rsd5f

      sum
      =====
      71680
```

- Version 15.0 and later:

```
1> select a.dbid, a.size, b.phyname
2> from sysusages a, sysdevices b
3> where b.name = "testdev"
4> and a.vdevno = b.vdevno
5> compute sum(size)
6> go
```

Subtract the sum from the total space on your physical device to determine the available space on the database device.

Additional information

Refer to “model database” in the *System Administration Guide: Volume 1* and “create database” in the *Reference Manual: Commands* for more information.

Version in which this error is raised

All versions

Error 1808

Severity

21

Message text

```
Crdb_disk: Getnext SCAN_NOINDEX on
sysdevices.status=DEFAULT failed to find default rows
```

Explanation This error occurs during execution of a `create` or `alter database` command when you do not specify a device and no disk device is defined as the default disk for an Adaptive Server.

By default, the master device is defined as a default disk. If you have used the stored procedure `sp_diskdefault` to turn off default for the master device:

```
1> sp_diskdefault master, defaultoff
2> go
```

and you have not used `sp_diskdefault` to specify another default disk device:

```
1> sp_diskdefault <dev_name>, defaulton
2> go
```

Where `<dev_name>` is the new default disk device. Afterward, you will not have any devices defined as default devices and error 1808 will occur when you run `create` or `alter database`.

Action Select one of the following options:

- Always use the `on <dev_name>` clause with `create` and `alter database` statements.
- Follow these steps:
 - a Use one of the methods below to determine whether you have any disk devices defined as default.
 - b If you do not have any disk devices defined as default, use `sp_diskdefault` to specify a device as a default disk device.

Method 1

Type:

```
1> sp_helpdevice
2> go
```

If you see a line such as:

device_name	physical_name	description
master	d_master	special,default disk,physical disk,17.00MB

where “default disk” is listed, that device will be used by `create` or `alter database` when the `on <dev_name>` clause is omitted for allocation for the database. If no entries have “default disk” in the description field, you do not have any disk devices defined as default.

Method 2

You can determine whether you have any disk devices defined as default by checking whether the first bit of the `status` column in `sysdevices` is turned on for any devices:

```
1> select status, name from sysdevices
2> where status & 1 = 1
3> go
```

```
status name
-----
3 master
```

Additional information	Refer to the <i>Reference Manual: Procedures</i> for information about <code>sp_diskdefault</code> .
Version in which this error is raised	All versions

Error 1809

Severity	14
Message text	CREATE DATABASE must be preceded by a 'USE master' command. Check with your DBO <or a user with System Administrator (SA) role> if you do not have permission to USE master.
Explanation	This error occurs when you try to create a new database without being in the master database.
Action	Before issuing a create database command, be sure you are in master: <pre>1> use master 2> go</pre> <p>If you do not have the necessary permissions to create a database, ask your Sybase System Administrator to grant you permission to use the <code>create database</code> command.</p>
Version in which this error is raised	All versions

Error 1810

Severity	16
Message text	<code>CREATE DATABASE failed because of incorrect size parameter(s). Total number of megabytes specified must be at least %d megabytes so that the Model Database can be copied to the new database.</code>
Explanation	When a create database command is issued, Adaptive Server makes a copy of the model database, which contains the system tables needed by each database. Error 1810 occurs when you try to create a database of a size smaller than the size of the model database. The default size of model is 2 or 3MB, depending on the version of Adaptive Server Enterprise you are running.
Action	<p>Make sure the size you plan to use in your create database command is at least as large as the model database (combine the data space size and the log space size).</p> <p>To check the size of model, use this command:</p> <pre>1> sp_helpdb model 2> go</pre>
Additional information	Refer to “model database” in the <i>System Administration Guide: Volume 1</i> and “create database” in the <i>Reference Manual: Commands</i> for more information.
Version in which this error is raised	All versions

Error 1813

Severity	16
Message text	<code>Cannot open new database '%.*s'. CREATE DATABASE is aborted.</code>
Explanation	<p>During a create database command, Adaptive Server opens the new database at several stages, first after inserting the rows into the system catalogs and later to initialize the log segment. Error 1813 occurs when Adaptive Server is unable to open the new database or initialize the log segment for the new database during a create database command.</p> <p>This error can be caused by:</p>

	<ul style="list-style-type: none"> • Uncleared allocation buffers. This can occur after a database has been dropped and then a <code>create database</code> command is issued. • Corrupt system catalogs such as <code>sysdatabases</code>. • An Adaptive Server problem.
Action	<p>If error 1813 occurs just after a database has been dropped, shut down and restart Adaptive Server to clear the allocation buffers and then re-enter your <code>create database</code> command.</p> <p>If this does not clear the problem, or you have not just dropped a database, call Sybase Technical Support.</p>
Additional information	<p>Have the information listed in “Reporting errors” on page 4 available when you call Sybase Technical Support, including the <code>dbcc checkcatalog</code> output for master.</p>
Version in which this error is raised	All versions

Error 1820

Severity	20
Message text	<p>This command adds <code>%S_MSG</code> space to disk '<code>%.s</code>', which previously contained only <code>%S_MSG</code>. You must specify <code>WITH OVERRIDE</code> to force this allocation.</p>
Explanation	<p>When a database is being extended or a new database created, Adaptive Server allocates space on the specified device (if no device is named, it uses the default device).</p> <p>Error 1820 is raised when allocating space for the database converts the device from a dedicated device that stores only data or log into a mixed-use device that contains both data and log on the same segment.</p> <p>The error may be detected:</p> <ul style="list-style-type: none"> • When you use the <code>alter database</code> or <code>create database</code> commands • When proxy databases are being created for the high availability (HA) feature (version 12.0 and later). <p>Error 1820 is raised with the following states:</p>

State	Meaning
1	Cannot allocate log space on a device dedicated to data. Device was named in the <code>create database/alter database</code> command.
2	Cannot allocate data space on a device dedicated to the log. Device was named in the <code>create database/alter database</code> command.
3	Cannot allocate log space on a device dedicated to data. Device was not explicitly specified. This state is also raised during HA proxy database creation (version 12.0 and later).
4	Cannot allocate data space on a device dedicated to the log. Device was not explicitly specified. This state is also raised during HA proxy database creation (version 12.0 and higher).

Action	Contact Sybase Technical Support with the information on hand that is listed in “Reporting errors” on page 4.
Additional information	For details about the management and recovery of the transaction log, see “Placing the Transaction Log on a Separate Device” in the <i>System Administration Guide: Volume 2</i> .
Version in which this error is raised	Version 11.0.3 and later

Error 1902

Severity	16
Message text	Cannot create more than one clustered index on table '%.*s'. Drop the existing clustered index '%.*s' before creating another
Explanation	<p>Adaptive Server uses a clustered index to sort rows so that their physical order is the same as their logical (indexed) order. The bottom or leaf level of a clustered index contains the actual data pages of the table.</p> <p>In a nonclustered index, the physical order of the rows is not the same as the indexed order. The leaf level of a nonclustered index contains pointers to rows on data pages.</p>

Many nonclustered indexes are allowed on a table, but only one clustered index per table is allowed. Error 1902 occurs when you try to create a second clustered index on the table named in the error message.

Note Using the primary key clause in a create table statement creates a unique clustered index.

Action	<ol style="list-style-type: none"> 1 List the existing indexes on the table: <pre>1> use <database_name> 2> go 1> sp_helpindex <object_name> 2> go</pre> <p>Where <i><database_name></i> is the database where the table named in the error message resides and <i><object_name></i> is the name of the table in the error message. The <i>index_description</i> column of the output shows whether the index is clustered or nonclustered.</p> 2 If a clustered index already exists, drop it. 3 Create a new clustered index.
Additional information	See “Using drop index and create clustered index” in the most recent version of <i>Performance and Tuning: Basics</i> for more information about drop index and create index.
Version in which this error is raised	All versions

Error 1903

Severity	16
Message text	%d is the maximum allowable size of an index. Composite index specified is %d bytes.
Explanation	<p>A composite index is any index that uses from 2 to 16 columns. Error 1903 occurs when the sum of the lengths of all the columns used in a composite index exceeds the allowable limits.</p> <p>Following is an example of a composite index using objects from the pubs2 sample database:</p>

```
1> create index example
2> on authors (phone, state, postalcode)
3> go
```

phone is defined as char(12), state as char(2), and postalcode as char(10). Since each char takes 1 byte of storage, the total length of this composite index is 24 bytes.

Action	Examine the lengths of the columns in your composite index and make sure that the lengths total no more than 600 bytes. Refer to “Using system-supplied datatypes” and “User-defined datatypes” in the <i>Transact-SQL User's Guide</i> for information about how to calculate the storage size for different Adaptive Server datatypes.
Additional information	Refer to “create index” in the <i>Reference Manual: Commands</i> for information about indexes.
Version in which this error is raised	All versions

Error 1916

Severity	16
Message text	CREATE INDEX options %s and %s are mutually exclusive.
Explanation	This error occurs when you try to create an index with two mutually exclusive create index options. When error 1916 occurs, no index is created.

The following options are not compatible with each other:

- ignore_dup_row
- allow_dup_row
- ignore_dup_key

The following table shows the options that can be used with each type of indexes:

Index type	Index option		
	ignore_dup_key	ignore_dup_row	allow_dup_row
Nonunique Nonclustered	Do not use	Do not use	Do not use
Nonunique Clustered	Do not use	Okay to use	Okay to use

Index option			
Unique Nonclustered	Okay to use	Do not use	Do not use
Unique Clustered	Okay to use	Do not use	Do not use

Action	<ul style="list-style-type: none"> • Use only one of the options. • Use the option appropriate for the type of index being created.
Additional information	Refer to “create index” in the <i>Reference Manual: Commands</i> .
Version in which this error is raised	All versions

Error 1928

Severity	16
Message text	<p>Version earlier than 15.0:</p> <pre>Cannot create clustered index on '%.*s' because it is partitioned.</pre> <p>Version 15.0 and later:</p> <pre>Cannot create clustered index '%.*s' on '%.*s' which has %d partitions because the maximum number of partitions allowed on a table with clustered index is %d.</pre>
Explanation	<p>By default, Adaptive Server stores a heap table’s data in one doubly linked chain of database pages. Adaptive Server inserts all new rows into the last page of the chain. A transaction holds an exclusive lock on the last page while inserting new rows, which can block other, concurrent transactions from inserting into the table.</p> <p>The <code>partition</code> clause of the <code>alter table</code> command allows you to partition user tables that do not have a clustered index. Partitioning creates additional page chains on the table, each with its own last page. This reduces page contention for concurrent inserts, and can also reduce I/O contention if the table exists on user-defined segments and is distributed over multiple physical devices.</p> <p>Since clustered indexes are not allowed on partitioned tables, error 1928 occurs when you try to create a clustered index on a partitioned table.</p>

Action	<p>If you want to create a clustered index on the table, use the unpartition clause of the alter table command to concatenate all partitions:</p> <pre>1> alter table <table_name> unpartition 2> go</pre> <p>Then create your clustered index.</p>
Additional information	Refer to the <i>Reference Manual: Commands</i> for information about the alter table command.
Version in which this error is raised	All versions

Error 2110

Severity	20
Message text	<pre>The proccreate() function is trying to put too many lock requests in the lock_requests[] array.</pre>
Explanation	<p>When you create and modify triggers (and other stored objects such as procedures, views, and rules), Adaptive Server uses the proccreate function to update the appropriate system tables. As an example, Adaptive Server uses this function to write the text of a trigger into the syscomments table. Depending on the action requested, proccreate must acquire locks on the objects being created and any objects affected by the created object.</p> <p>For example, since a table can only have one trigger for each operation (create, update or delete), a new trigger on a table for the same operation overwrites any previous trigger for that operation. proccreate must be able to lock the new trigger as well as the existing trigger.</p> <p>Error 2110 is raised when you attempt to create a new trigger or other stored object on a table, or when you attempt to overwrite an existing trigger, and proccreate is unable to acquire the necessary locks on the objects in question.</p>

State	Meaning
1	Could not acquire a lock for a trigger or other object being created.
2	Could not lock an existing delete trigger on this object.
3	Could not lock an existing insert trigger on this object.
4	Could not lock an existing update trigger on this object.
5	Could not lock the table to which the trigger is being attached.

Action	<p>Error 2110 breaks the connection to the Adaptive Server.</p> <p>Take the following steps to correct the problem.</p> <ol style="list-style-type: none"> 1 Retry the object creation. 2 If the error was raised with state 5, and the target table is newly created, check your trigger creation code to be sure that the table's <code>create</code> statement is committed before attempting trigger creation. 3 If the above steps do not resolve the problem, error 2110 may be due to an Adaptive Server problem. Contact Sybase Technical Support for assistance.
Additional information	When calling Technical support, have the information available that is listed in "Reporting errors" on page 4, including the text of the trigger that raises the error.
Version in which this error is raised	All versions

Error 2714

Severity	16
Message text	There is already an object named '%.*s' in the database.
Explanation	<p>This error occurs when Adaptive Server tries to create an object that already exists.</p> <p>Error 2714 can occur:</p> <ul style="list-style-type: none"> • When you are creating procedures, triggers, tables, views, or rules, if you specify an object that already exists. • If you changed the sort order to be case-insensitive, then tried to create an object that is spelled the same way as an existing object with a different case (for example, PROC1 and proc1). <p>When using temporary tables:</p> <ul style="list-style-type: none"> • Version pre-15.0, the first 13 bytes of the table name, including the pound sign, must be unique for the current session. Long temporary table names are truncated to 13 characters with the pound sign.

- Versions 15.0 and later, if you create a temporary table with a name requiring fewer than 238 bytes, Adaptive Server adds a 17-byte suffix to ensure that table name is unique. If you create a temporary table with a name of more than 238 bytes, Adaptive Server uses only the first 238 bytes (including pound sign), then adds the 17-byte suffix.
- An attempt to create a second temporary table with the same name in the same session will raise error 2714. (In contrast, temporary tables created by specifying `database_name.owner.table_name` may be shared among sessions, not only within one session.)

Action	Make sure the name of the object you are trying to create does not already exist. If you have stored procedures that create objects, make sure there is a check for the existence of the object in the stored procedure (before it is created). If you want to re-create an object within a stored procedure, drop the object before you re-create it.
Additional information	Refer to the <i>Reference Manual: Commands</i> for information about creating and dropping objects.
Version in which this error is raised	All versions

Error 2729

Severity	16
Message text	Procedure '%.*s' group number 1 already exists in the database. Choose another procedure name or procedure number.
Explanation	<p>Procedures that are meant to be used together in an application can be created with the same name and different version numbers, using an option of the create procedure statement. For example, procedures used with the application ticket can be named <code>ticketproc;1</code>, <code>ticketproc;2</code>, and so on. The numbers are referred to as the group numbers of the procedure.</p> <p>Error 2729 occurs when you attempt to:</p> <ul style="list-style-type: none">• Create a procedure without specifying any group number, and a procedure by that name already exists. Adaptive Server assumes that you want to create a procedure in the same group as the existing procedure, but the command fails since no group number is specified, and the default group number 1 already exists.

Action	<ul style="list-style-type: none"> • Re-use an existing group number when creating procedures with the same name. <p>To resolve the problem, do one of the following:</p> <ul style="list-style-type: none"> • Supply a different procedure name or • Supply the same procedure name, but include a unique procedure group number. For example: <pre>1> create procedure myproc;2 ...</pre> <p>This results in two procedures with the same name and group numbers 1 and 2, respectively.</p>
Version in which this error is raised	All versions

Error 2753

Severity	20
Message text	Received an error code of '%d' from 'intl_strblist()'.
Explanation	<p><code>intl_strblist</code> is an Adaptive Server function used to convert a message from the server's default language to the client's language. Error 2753 is raised when there is an inconsistency in language between the client session and the server's default language, and <code>intl_strblist</code> is unable to convert the message to the client format. The error code denotes the type of conversion problem that occurred.</p> <p>For example, this error may be raised when you invoke an isql session with the <code>-J</code> option to specify a client character set different from <code>us_english</code> and run certain <code>dbcc</code> commands. The error may occur even if Adaptive Server is configured to use the alternate language.</p> <p>This is not a serious error and does not cause any problems in your database. Depending on the type of error, Adaptive Server may be able to print the message string in <code>us_english</code>.</p>
Action	<p>Use one of the following actions to address the 2753 error:</p> <ul style="list-style-type: none"> • Invoke the client session without the <code>-J</code> option. • Use <code>set language us_english</code> before operations that could produce error 2753.

Version in which this error is raised All versions

Error 2762

Severity 16

Message text The '%s' command is not allowed within a multi-statement transaction in the '%.*s' database.

Explanation SQL commands are grouped into these categories:

- SQL commands that are not allowed in transactions at all.
- SQL commands, such as data definition language (DDL) commands, that are allowed in transactions only if the required database option (`ddl in tran`) is set to “true”.
- SQL commands that are run across databases to create, alter or drop objects in another database, and are allowed in transactions only if the required database option (`ddl in tran`) is set to “true” for that database.

Error 2762 typically occurs in the context of data definition commands such as creating, altering, or dropping objects. It occurs when Adaptive Server detects a command that is not allowed in a multi-statement transaction for the specified database. A multi-statement transaction is a set of commands prefaced with the `begin transaction` command, or when chained mode is on.

Error 226 is a similar error that may be raised along with error 2762. While Error 2762 typically involves commands that are run across databases, error 226 is raised when the command affects only the local database.

The following commands are never allowed in multi-statement transactions:

- `alter database`
- `create database`
- `dbcc reindex`, `dbcc fix_text`
- `disk init`
- `drop database`
- `dump database`, `dump transaction`
- `load database`, `load transaction`

- reconfigure
- select into
- set transaction isolation level
- truncate table
- update statistics

The following DDL commands are not normally allowed in multi-statement transactions but you can use them if you use `sp_dboption` to set `ddl in tran` to “true” first:

- create default, create index, create procedure, create rule, create schema, create table, create trigger, create view
- drop default, drop index, drop procedure, drop rule, drop table, drop trigger, drop view
- grant
- revoke

Action

The literal ``%. *s'` in the error message is the name of the database specified in the command. Note the database name before choosing any corrective action.

For commands involving *tempdb* database

Error 2762 can be raised when a command creates or drops temporary objects (objects in *tempdb*) within the context of a multi-statement transaction. For example, this code may generate a 2762 error:

```
1> begin transaction
2> create table #cities
3> (city_name char(15) not null)
4> commit transaction
5> go
```

The error is raised when `ddl in tran` is set to “false” in *tempdb*. Use one of the following strategies to correct this error:

- Use a permanent object name, so that *tempdb* is not affected.
- Execute the command outside a multi-statement transaction.

Warning! Using data definition language commands on *tempdb* within transactions may cause concurrency problems in *tempdb*. Always leave `ddl in tran` set to “false” in *tempdb*.

For commands not involving *tempdb* database

If the DDL command is allowed in a multi-statement transaction when `ddl in tran` is set to “true”, set `ddl in tran` to “true” before running the transaction. You can check the current setting of `ddl in tran` with `sp_helpdb`.

Warning! DDL commands hold locks on system tables such as `sysobjects` and this can affect performance. Avoid using them inside transactions; if you must use them, keep the transactions short.

If the command is never allowed in a multi-statement transaction, execute it outside the multi-statement transaction.

Additional information

Some applications take SQL statements as input and run them automatically. If the application uses `begin` and `commit` or `rollback transaction` to surround those statements, error 2762 may occur. Refer to the user guide for your application to determine if this is the case.

Version in which this error is raised

All versions

Error 12818

Severity

16

Message text

Cannot create an object of this type in system-created proxy database.

Explanation

Adaptive Server Enterprise’s high availability (HA) support allows two nodes in a cluster to take over each other’s workload in the event of a failure. The server that fails is called the primary, and the server that takes over the workload is the secondary. Movement of data access services from a failed primary server to a companion secondary server is known as failover.

In the HA architecture, the secondary companion is configured with proxy databases which act as placeholders; in the event of failover, proxy databases are dropped and replaced with databases mounted and recovered from the primary server.

Error 12818 is raised when you attempt to execute certain data definition language (DDL) commands in a proxy database, including:

- `create/drop procedure`

- create/drop view
- create/drop trigger

These commands are not available in a proxy database, which is a system database reserved solely for HA use in the event of failover.

Action	No action is needed. User objects should not be placed in a proxy database.
Additional information	For details about the DDL commands supported in primary and secondary servers in an HA environment, see <i>Using Sybase Failover in a High Availability System</i> .
Version in which this error is raised	Version 12.0 and later

Error 12842

Severity	18
Message text	Adaptive Server cannot process this %s statement due to one or more errors. Please retry this command if it failed due to a deadlock. Please refer to your application log or Adaptive Server errorlog for more information on the cause of this failure.
Explanation	Adaptive Server internally creates a hidden materialized computed column when a functional index is built. When the functional index is dropped, the corresponding hidden computed column also should be dropped. This error indicates that an error occurred while trying to drop the hidden computed column as part of dropping the functional index.
Action	If there is a deadlock, rerun your command. If no deadlock was indicated, this may indicate an internal Adaptive Server error.
Additional information	A configurable shared memory dump (CSMD) on error 12842 may assist in determining root cause. Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 12881

Severity	17
Message text	<p>There is insufficient memory to allocate the structures needed to set local switches.</p> <p>Contact your System Administrator</p>
Explanation	<p>Starting with Adaptive Server versions 12.5.4 and 15.0.2, trace flags can be turned on and off using the <code>set switch</code> command (rather than <code>dbcc traceon</code>).</p> <p>For each user process and some system processes, Adaptive Server maintains a process status structure (PSS) to track the state of processes and resource use. Memory is dynamically allocated from the PSS session pool to turn local switches on, when needed. This memory is freed if the last local switch is turned off. The PSS session pool is created at start time and can be used by all processes. Processes allocate memory from this pool as and when needed. Error 12881 is raised when an attempt to allocate memory from the PSS session pool fails.</p>
Action	<p>Try setting this local switch later when more memory becomes available or consider increasing heap memory per user. See the chapter “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 2</i>.</p>
Version in which this error is raised	Versions 12.5.4 and 15.0.2 or later

This chapter contains error message write-ups for Adaptive Server Enterprise's high availability support.

Error	Page
Error 2243	263

Error 2243

Severity	16
Message text	<code>QUIESCE DATABASE failed with error %ld. See the error log for more detail.</code>
Explanation	<p>Popular techniques for backing up an Adaptive Server database, without the need for a full dump and load cycle, include unlinking a mirror device and using file system backup utilities outside the server. In version 12.0, <code>quiesce database</code> allows you to unlink mirrors or back up devices at the operating system level without having to first shut down the server. <code>quiesce database hold</code> establishes a quiesce point and prevents updates in the specified databases. <code>quiesce database release</code> allows transactions to resume once the device is unlinked or the external backup completes. <code>quiesce database hold</code> creates a child process, known as a quiesce db agent, to carry out the tasks necessary to establish the quiesce point.</p> <p>Error 2243 is raised when you issue <code>quiesce database hold</code>, but the quiesce db agent detects a condition which prevents it from creating the quiesce point. Possible causes include:</p> <ul style="list-style-type: none"> • <code>dump database</code> was not executed at least once on the database prior to issuing <code>quiesce database hold</code> • <code>dump transaction</code> has been executed on the database since the last <code>dump database</code>; • The timeout limit was reached while the agent waited for transactions in prepare state to finish.

Error 2243

Action	Note the error number specified in the 2243 message text. Locate this error in the server error log to determine the cause of the problem. Take the appropriate action, such as issuing a dump database, to resolve the error.
Version in which this error is raised	Version 12.0 and later

Character Set Conversion Errors (2400s)

This chapter contains error messages for Adaptive Server Enterprise's character set conversion that is used when the client uses a different character set than Adaptive Server.

Error	Page
Error 2401	265
Error 2402	266
Error 2409	268

Error 2401

Severity	11
Message text	Character set conversion is not available between client character set '%.*s' and server character set '%.*s'.
Explanation	<p>Character set conversion is the change of the encoding scheme of a set of characters on the way into or out of Adaptive Server. Conversion is used when Adaptive Server and a client communicating with it use different character sets. For example, if Adaptive Server uses character set ISO 8859-1 (ISO_1) and a client uses Code Page 850, character set conversion must be turned on so that both Adaptive Server and the client interpret data passing back and forth in the same way.</p> <p>Error 2401 occurs when an invalid character set conversion is requested (for example, ASCII_8 with ISO 8859-1). When this error is raised, no conversion is performed.</p> <p>Error 2401 can occur:</p>

- When a client requests a connection, Adaptive Server checks whether it can convert from the client's character set to its own character set. If it cannot do the requested conversion because conversion is not available between the two character sets, this error is raised.
- If you use the `set char_convert` command to start conversion between the Adaptive Server character set and a different character set and conversion for those character sets is not available, this error is raised.

Action	Refer to the table "Supported languages and character sets" in the chapter "Configuring Character Sets, Sort Orders, and Languages" in the <i>System Administration Guide: Volume 1</i> and check the table of valid conversion matches. Make sure you are not requesting conversions involving incompatible character sets. Consider using Unicode conversions to convert between incompatible character sets.
Additional information	Character set conversion is only available for certain character sets. Valid conversion matches are listed in the <i>System Administration Guide: Volume 1</i> in the chapter "Configuring Client/Server Character Set Conversions."
Version in which this error is raised	All versions

Error 2402

Severity	16
Message text	Error converting client characters into server's character set. Some character(s) could not be converted.
Explanation	<p>This error occurs during insertion of data (<code>insert</code> or <code>bcp</code>) when Adaptive Server fails to convert a character to the required character set.</p> <p>Error 2402 usually occurs when:</p> <ul style="list-style-type: none">• The character exists in the client character set but it does not exist in the Adaptive Server character set.• The character exists in both the client and the Adaptive Server character set, but is represented by a different number of bytes in the client character set than in the Adaptive Server character set. <p>This error occurs during normal processing and it prevents query execution.</p>
Action	The following options are available for recovering from error 2402.

Changing the data

Modify the incoming data so that it contains characters recognizable by Adaptive Server.

Turning off character set conversion

If the error occurs while you are using `isql`, `bcp`, or `defncopy`, you can use the `-J` (UNIX and Windows) command-line option with no character set name to set the client's character set to NULL. If you use this command-line option without specifying a character set name, no conversion takes place and no error message is sent. As a result, some characters sent by the client to the Adaptive Server may not be interpreted correctly by the Adaptive Server and vice versa. (If only 7-bit characters are being handled, no incorrect interpretation will take place.)

Otherwise, you can turn off conversion so that characters are sent and received unchanged with the following command:

```
1> set char_convert off
2> go
```

Turning off character set conversion error reporting

You can turn off the printing of error messages with the following command:

```
1> set char_convert on with no_error
2> go
```

Bytes that cannot be converted are replaced with an ASCII question mark (“?”).

Note The preceding commands turn off error reporting only when characters are being sent from Adaptive Server to the client. Error 2402 is still raised if a client sends a character that Adaptive Server cannot convert.

Additional information

Refer to “Character set conversion in Adaptive Server” in chapter 10 of the *System Administration Guide: Volume 1* for details about character set conversion.

Version in which this error is raised

All versions

Error 2409

Severity	11
Message text	Cannot find the requested character set in Syscharsets: name = '%.*s'.
Explanation	<p>Information about the character sets that are defined for use by Adaptive Server is stored in <code>syscharsets</code>. Error 2409 occurs when Adaptive Server cannot find a character set needed for character set conversion.</p> <p>Error 2409 can occur:</p> <ul style="list-style-type: none">• When a client requests a connection, Adaptive Server checks whether it can convert from the client's character set to its own character set. If Adaptive Server cannot do the requested conversion because a character set is not found, error 2409 is raised.• If you use the <code>set char_convert</code> command to start conversion between the Adaptive Server character set and a different character set and the character set you specify does not exist in <code>syscharsets</code>, error 2409 is raised.
Action	<p>Make sure all necessary character sets are loaded, including the client's character set (as shown in the error message output):</p>

```
1> use master
2> go

1> select csid, id, name from syscharsets
2> go
csid id  name
---- -
0    0  ascii_8
0    1  iso_1
1   50 bin_iso_1
```

If you are missing needed character sets, follow the directions in the Adaptive Server installation and configuration guide to load them using the `sybinit` utility.

Standalone utilities

By default, for standalone utilities, the character set used for the client is the default for the platform. This may not be the correct character set. If the 2409 error is occurring while you are using `isql`, `bcp`, or `defncopy`, use the `-J` (UNIX and Windows) command-line option to specify the client character set.

Open Client applications

For Open Client applications such as DB-Library, you can specify the character set for the client through the login packet. Refer to the reference manual for the application and *Open Server 15.0*, *Open Client 15.0* and *SDK 15.0* for details.

Additional information

Refer to the chapter “Configuring Client/Server Character Set Conversions” in the *System Administration Guide: Volume 1*.

Version in which this error is raised

All versions

***dbcc* Errors (2500s, 7900s, 12000s, 15000s)**

This section contains error messages for *dbcc* (database consistency checker) commands.

Note Messages that ordinarily have severities greater than 16 will show severity 16 when raised by *dbcc checktable* and *dbcc checkalloc* so that checks continue with the next object.

Error	Page
Error 2501	272
Error 2502	273
Error 2503	275
Error 2507	277
Error 2509	279
Error 2510	281
Error 2511	283
Error 2513	284
Error 2514	285
Error 2517	287
Error 2520	289
Error 2521	290
Error 2524	294
Error 2525	296
Error 2529	299
Error 2540	301
Error 2546	305
Error 2547	310
Error 2550	312
Error 2558	314
Error 2559	321
Error 2572	324

Error	Page
Error 2573	325
Error 2574	325
Error 2575	327
Error 2578	328
Error 2582	330
Error 2583	332
Error 2591	334
Error 2596	335
Error 7901	338
Error 7902	340
Error 7928	343
Error 7930	343
Error 7939	346
Error 7940	348
Error 7949	350
Error 7989	352
Error 9961	353
Error 12964	354
Error 15056	355
Error 15080	356
Error 15082	356
Error 15084	357

Error 2501

Severity 16

Message text Table named %.*s not found; check sysobjects.

Explanation This error occurs when a dbcc command is run if Adaptive Server is unable to find the table name (or ID) in sysobjects.

Error 2501 can occur when running any of the following dbcc commands:

- dbcc checktable
- dbcc checkdb
- dbcc tablealloc

- dbcc reindex
- dbcc fix_text

Warning! Because this error can be serious and indicates possible corruption, you must correct the error immediately.

Action If this error occurs on a command for which you have specified a table name, check to make sure you have spelled the table name correctly:

```
1> select * from sysobjects
2> go
```

If this error occurs because the entry for an object has been removed from `sysobjects` but the object is still referenced by some other system tables, call Sybase Technical Support. They will help you try to remove the references to that object from other system tables.

Additional information Before calling Technical Support, have the information available that is listed on “Reporting errors” on page 4, including the of `dbcc checkalloc` and `dbcc checkdb`.

Version in which this error is raised All versions

Error 2502

Severity 16

Message text Table Corrupt: A page is linked in more than one chain;
check this page:page number=%ld allocation status=%d

Note This error can be caused by a hardware problem.

Explanation This error occurs when Adaptive Server detects an inconsistency in the page chains associated with an object. For example, a page may be connected in more than one chain. If this error is detected by a `dbcc` command, error 605 can also occur.

Warning! If error 605 occurs with this error, use the instructions under “Error 605” on page 72 to remedy that error or call Sybase Technical Support.

Action

This is a serious error that must be corrected immediately.

- 1 Note the page number that appears in the error message. Use the procedure “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide to determine the table associated with the corrupt page chain.

Warning! The procedure referenced in step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 Depending on the type of object, select the most appropriate procedure:
 - If the object with the error is not a system table (a system table’s object ID is less than 100), go to step 3.
 - If the object with the error is a system table and the index ID is not 0 (zero), see “Fixing a system table corrupted index” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery* for instructions on how to repair the system table index.
 - If the index ID is 0 (zero) , contact Sybase Technical Support. They may be able to help you repair the corruption, but it may be necessary to restore from clean backups.
- 3 If the object with the error is not a system table, use one of the following procedures below to correct the problem, depending on the kind of object is involved:
 - If the page belongs to an index, drop and re-create the index.
 - If the page does not belong to an index, recover your database from a known clean backup or recover the tables not affected by this error using the `bcp` utility. See “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery* for instructions.
- 4 Examine your operating system error log and the Adaptive Server error log to determine if hardware errors caused the corruption.

Additional information	Before calling Sybase Technical Support, have information available that is listed in “Reporting errors” on page 4, including the output of the dbcc page command.
Version in which this error is raised	All versions

Error 2503

Severity 16

Message text **Versions earlier than 15.0:**
Table Corrupt: Page linkage is not consistent; check the following pages: (current page#=%ld; page# pointing to this page=%ld; previous page# indicated in this page=%ld)

Version 15.0 and later:

Table Corrupt: Page linkage is not consistent. Check the following pages: current page number = %d, partition ID = %d; page number of the page pointing to this page = %d; previous page number indicated in this page = %d

Note This error can be caused by a hardware problem.

Explanation This error occurs when Adaptive Server detects an inconsistency in the page linkage of one of the page chains associated with a table. There is one doubly-linked page chain for the table’s data as well as one for each index level. For example, a page’s “next” pointer points to a page whose “previous” pointer points to a different page.

Action This is a serious error that must be corrected immediately.

- 1 Note the first page number in the error message. Use the procedure “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide to determine the object associated with the “current page”.

Warning! The procedure referenced in step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 Depending on the type of object, select the appropriate procedure:
 - If the object with the error is not a system table (a system table’s object ID is less than 100), go to step 3.
 - If the object with the error is a system table and the index ID is not 0 (zero), see “Fixing a system table corrupted index” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery* for instructions on how to repair the system table index.
 - If the index ID is 0 (zero), contact Sybase Technical Support to help you repair the corruption, but it may be necessary to restore from clean backups.
- 3 If the object is not a system table, perform one of the procedures below, depending upon the type of object involved in the error.

Current page belongs to a nonclustered index

If the “current page” belongs to a nonclustered index (index ID is between 2 and 250), run `dbcc checktable` to verify that the data page linkage is good. If `checktable` runs cleanly, drop and re-create the index.

Current page belongs to a clustered index

If the “current page” belongs to a clustered index, drop all the indexes on that table, re-create the clustered index first, then re-create the remaining nonclustered indexes.

Current page does not belong to an index

If the “current page” does not belong to an index at all (index ID is 0 (zero) or 255) and there were no errors other than error 2503, perform these steps:

- 1 Use `bcpl` or `select into` to retrieve the data from the affected table. See “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery* for instructions.
- 2 If errors still exist, restore from clean backups or call Sybase Technical Support.
- 3 Examine your operating system error log and the Adaptive Server error log to determine if hardware errors may have caused the corruption.

Additional information

For information on locating device fragments, refer to “Determining the physical device on which a database resides” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery*.

Version in which this error is raised

All versions

Error 2507

Severity

16

Message text

Versions earlier than 15.0:

```
Table Corrupt: Offset table is incorrect (page#=%ld
row#=%d offset in offset table is=%d correct offset
is=%d)
```

Version 15.0 and later:

```
Table Corrupt: Offset table is incorrect (page number
%d, partition ID %d; row number %d; offset in the offset
table is %d; the correct offset is %d).
```

Note This error can be caused by a hardware problem.

Explanation

An “offset table” is stored at the end of every data page and indicates where rows are located on that page. Adaptive Server uses the offset table to quickly access the rows in each page.

Error 2507 occurs when `dbcc checkdb` or `dbcc checktable` detects inconsistencies in the offset table of a row on the page specified by the error message.

Action

Refer to “Identifying and fixing a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery*.

Warning! Some data on this page might be lost if you recover your table using `bcp` or `select into` (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

Before dumping your database, make sure it is working correctly. More specifically, run the following commands prior to each dump:

- 1 `dbcc checkdb`.
- 2 `dbcc checkalloc` or `dbcc checkalloc` with the `fix` option. Refer to “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery* for information about running these commands in multi-user mode and preventing false allocation errors from `dbcc` commands.

Note Alternatively, using `dbcc checkstorage` with `checkverify` can examine for most types of consistency problems in less time.

Hardware errors

In addition to the preceding causes, hardware errors on database devices can cause error 2507. To rule out this possibility:

- Check the Adaptive Server error log to determine whether there are other indications of hardware problems, such as kernel messages reporting I/O errors.
- Check the hardware error log or diagnostics utilities for I/O errors.

If problems persist, refer to the error documentation in this manual or contact Sybase Technical Support.

Additional information

Refer to `bcp` in the most recent version of the Adaptive Server Enterprise *Utility Guide*.

Refer to “select” in the *Reference Manual: Commands* for information about the select into command.

Before calling Sybase Technical Support, have information available that is listed in “Reporting errors” on page 4, including the output from `dbcc checkdb`, `dbcc checktable`, and `dbcc checkalloc`.

Version in which this error is raised

All versions

Error 2509

Severity

16

Message text

Versions earlier than 15.0:

Table Corrupt: The row number and offset of each row in the page should have a matching entry in row number table; check this page (page#=%d row#=%d offset in row number table=%d)

Version 15.0 and later:

Table Corrupt: The row number and offset of each row in the page should have a matching entry in the row number table. Check page number %d, partition ID %d (row number %d, offset in row number table is %d)

Note This error can be caused by a hardware problem.

Explanation

The offsets for data or index rows are stored at the end of every page (in the row number table) and indicate where a certain row is located on that page.

Error 2509 occurs when the `dbcc checkdb` or `dbcc checktable` command detects that a row does not have an entry matching its offset (location) on the page in the row offset table. Attempts to delete the offending row will result in error 631 (see “Error 631” on page 91), while attempts to select the offending row may be successful.

This problem is usually a result of an Adaptive Server internal error, but may also be caused by:

- Hardware failure
- Sybase system administration problems

- UNIX system administration problems

Action

First, rule out any of the above-mentioned causes for this error by referring to the appropriate sections in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery*.

After eliminating more serious errors on the table, follow these steps to correct error 2509:

- 1 Use the procedure “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide to identify the table and index that correspond to the page number from the error message text.

Warning! The procedure referenced in step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 Depending on the type of object, select the most appropriate procedure:
 - If the object with the error is not a system table (object ID is more than 100), go step 3.
 - If the object with the error is a system table and the index ID is not 0 (zero), see “Fixing a system table corrupted index” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery* for instructions on how to repair the system table index.
 - If the index ID is 0 (zero), contact Sybase Technical Support to help repair the corruption; although it may be necessary to restore from clean backups.
- 3 If the object with the error is a user table, use one of these methods to clear the 2509 error:

- Create a clustered index on the corrupted table. Creating a clustered index copies the entire table onto new data pages, and overwrites the row number table on each page. If a clustered index already exists on the table identified in step 1, drop the clustered index and re-create it.

Warning! If you have other serious errors on this table (errors 614, 2506, 2507, or 2524), eliminate these errors first, before you create a clustered index. Refer to the documentation in this manual for information on how to eliminate the other errors.

- Select the entire table into a new table, and drop the old table.
- Bulk copy the affected table out, drop and re-create the table, and bulk copy the table back in.

Warning! Some data on this page could be lost if you recover your table using `bcp` or `select into` (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

For more information about copying a table in a new table or file, refer to “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery* for instructions.

Additional information

Refer to “create index” and “select” in the *Reference Manual: Commands and in the Transact-SQL User's Guide*.

Refer to `bcp` in the most recent version of the Adaptive Server Enterprise *Utility Guide*.

Version in which this error is raised

All versions

Error 2510

Severity

16

Message text

Versions earlier than 15.0:

Key mismatch between index page %ld, at offset %d and

```
data page %ld (row ID %d) in database '%.*s'. Drop and
re-create index id %d of table '%.*s'.
```

Version 15.0 and later:

```
Key mismatch between index page %d, partition ID %d; at
offset %d; and data page %d (row ID %d), partition ID
%d. Drop and re-create index ID %d of table '%.*s' in
database '%.*s'.
```

Note This error can be caused by a hardware problem.

Explanation

This message indicates that an index is inconsistent with the table data that it represents.

Action

Use the this procedure to recover from the error:

- 1 Record the index page number and the data page number from the error text. Refer to “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide to identify the table and index that correspond to the data and index page number. Record the object ID and index ID.

Warning! The procedure referenced in step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 Depending on the type of object, select the most appropriate procedure:
 - If the object with the error is not a system table (its object ID is 100 or greater), go to step 3.
 - If the object with the error is a system table, see “Fixing a system table corrupted index” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery* for instructions on how to repair the system table index.
- 3 If the object with the error is a user table, find the name of the index involved by executing this query:

```
1> select name from sysindexes
2> where indid = <index_id>
3> and id = <object_id>
4> go
```

- 4 Drop and re-create the index. Refer to “drop index” and “create index” in the *Reference Manual: Commands* for more information.
- 5 Run `dbcc checktable` and `dbcc tablealloc` on the affected table to verify that all problems have been corrected at the table level.

If problems persist, contact Sybase Technical Support. Have the output from the appropriate `dbcc` commands available when you call.

Version in which this error is raised

All versions

Error 2511

Severity

16

Message text

Versions earlier than 15.0:

```
Keys of index id %d for table '%.*s' in %S_MSG not in
correct order. Drop and re-create the index. (index page
%d)
```

Version 15.0 and later

```
Keys of index ID %d for table '%.*s' in %S_MSG are not
in the correct order. Drop and re-create the index
(index page %d, partition ID %d).
```

Note This error can be caused by a hardware problem.

Explanation

This error occurs when an index is not correctly ordered.

There are three possible values of “%S_MSG” from the error message text:

- Data page (only on clustered indexes)
- Index page
- Leaf page

Action

- 1 Note the index ID and table name displayed in the error message output.

- 2 Depending on the type of object, select the most appropriate procedure:
 - If the object with the error is not a system table (object ID is more than 100), go to step 3.
 - If the object with the error is a system table, see “Fixing a system table corrupted index” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery* for instructions on how to repair the system table index.
- 3 If the object with the error is a user table, find the name of the index involved by executing this query:

```
1> select name from sysindexes
2> where indid = <index_ID>
3> and id = <object_ID>
4> go
```
- 4 Drop and re-create the index. Refer to “drop index” and “create index” in the *Reference Manual: Commands* for more information.
- 5 Verify that all problems have been resolved at the table level by running the dbcc checktable and dbcc tablealloc commands on the affected table.

If problems persist, contact Sybase Technical Support. Have the output from the appropriate dbcc commands available for their review.

Version in which this error is raised

All versions

Error 2513

Severity	16
Message text	Table Corrupt: Object id %ld (object name = %S_OBJID) does not match between %.*s and %.*s
Explanation	This error occurs when dbcc checkcatalog finds that an object exists in one system table but not in another system table where Adaptive Server expected to find it.

Warning! This is a serious error that indicates corruption.

Error 2513 occurs with the following states:

State	Meaning
1	The object was found in <code>syscolumns</code> but not in either <code>sysobjects</code> or <code>systypes</code> .
2	The object (a view) was found in <code>sysobjects</code> but not in <code>syscolumns</code> .
3	The object (a view) was found in <code>sysobjects</code> but not in <code>syscomments</code> .
4	The object (a view) was found in <code>sysobjects</code> but not in <code>sysprocedures</code> .
5	The object (a stored procedure) was found in <code>sysobjects</code> but not in <code>syscomments</code> .
6	The object (a stored procedure) was found in <code>sysobjects</code> but not in <code>sysprocedures</code> .
7	The object (a rule or default) was found in <code>sysobjects</code> but not in <code>syscomments</code> .
8	The object (a rule or default) was found in <code>sysobjects</code> but not in <code>sysprocedures</code> .
9	The object (a table) was found in <code>sysobjects</code> but not in <code>syscolumns</code> .
10	The object (a table) was found in <code>sysobjects</code> but not in <code>sysindexes</code> .
11	The object (a table) was found in <code>sysindexes</code> but not in <code>sysobjects</code> .

Action	Call Sybase Technical Support to help you delete the object that is causing the error. However, because other objects may reference that object, deleting it cleanly may be difficult. If this is the case, recovering from backups may be a better choice.
Version in which this error is raised	All versions

Error 2514

Severity	16
Message text	Table Corrupt: Type id %ld (type name = %.*s) does not match between %.*s and %.*s
Explanation	<code>syscolumns</code> contains a row for every column in a table, and <code>systypes</code> contains a row for every type in a table. Error 2514 occurs when <code>dbcc checkcatalog</code> detects a type mismatch between the <code>systypes</code> and the <code>syscolumns</code> system tables.

Error 2514 is most commonly caused by using the `select into` command across databases when at least one column from the original table is a user-defined datatype. It can also be caused by cross-database views and applications that map data across databases when a user-defined datatype involved in the operation has conflicting definitions in the two databases.

Note Error 2514 often prints out a `usertype` value that seems out of place: 0 (zero) or a large number. If your 2514 error does not match this description, call Sybase Technical Support for further assistance.

Action

Use the following procedure to find the incorrect entry in `syscolumns` and replace it with the correct value. Note that you actually type “c” and “t”; they are not variables. This method allows Adaptive Server to compare all tables to search for the mismatch.

- 1 Identify the rows containing incorrect entries:

```
1> select c.name, c.usertype, c.type
2> from syscolumns c
3> where not exists
4> (select * from systypes t
5> where t.usertype = c.usertype)
6> go
```

name	usertype	type
-----	-----	-----
PartNumber	0	56

More than one row may be returned.

- 2 Search for the correct value for each incorrect `syscolumns` `usertype` by querying the `systypes` `usertype` column, using the value reported in `type`. Do this for each row reported above in step 1:

```
1> select t.name, t.usertype, t.type
2> from systypes t
3> where t.type in ( 56 )
4> and t.usertype < 100
```

Results of the query should resemble this:

name	usertype	type
-----	-----	-----
int	7	56

The value returned under `usertype` is the correct value that should appear in the `usertype` column of `syscolumns`, “7” in this example.

3 Change the incorrect row value in the `systypes` column `usertype` with the correct value that you obtained in step 2:

a Enable updates to system tables (substitute your site's values for "7," "0," "56" and "PartNumber" in this example query):

```
1> sp_configure "allow updates", 1
2> go

1> begin transaction
2> update syscolumns set usertype = 7
3> where usertype = 0 and type = 56
4> and name = "PartNumber"
5> go
```

b To commit the transaction, type:

```
1> commit transaction
2> go
```

c Repeat the preceding steps (a – b) for any other incorrect rows, substituting the required values for `usertype`, `type`, and `name` in the example query.

d When you have updated all incorrect rows, disable updates to system tables:

```
1> sp_configure "allow updates", 0
2> go

1> checkpoint
2> go
```

4 Run `dbcc checkcatalog` again to verify that the problem is corrected.

Version in which this error is raised

All versions

Error 2517

Severity

16

Message text

Table Corrupt: Procedure id %ld (procedure name = %S_OBJID) does not match between %.*s and %.*s

Explanation	This error occurs when dbcc checkcatalog finds that an entry for a compiled object (for example, a stored procedure or trigger) exists in the table sysprocedures and that entry does not exist in the sysobjects table.
Action	<ol style="list-style-type: none">1 If the error occurred on an object in sysprocedures in master, enable updates to system tables:<pre>1> use master 2> go 1> sp_configure "allow updates", 1 2> go</pre>2 Determine how many rows contain the object entry:<pre>1> use <database_name> 2> go 1> select * from sysprocedures 2> where id = <procedure_ID> 3> go</pre><p>Where <i><database_name></i> is the name of the database that contains the table with the object and <i><procedure_ID></i> is the ID of the procedure in the 2517 error.</p>3 If you need to recover the stored procedure, retrieve the stored procedure's text:<pre>1> sp_helptext <stored_proc_name> 2> go</pre>4 Remove the object entry from sysprocedures:<pre>1> begin transaction 2> delete sysprocedures 3> where id = <object_ID> 4> go</pre><p>Where <i><object_ID></i> is the ID of the procedure in the 2517 error.</p>5 Check the results carefully, and commit the transaction only if the update affected the expected number of rows. If the update affected more than the expected number of rows, roll back the transaction. To commit the transaction, type:<pre>1> commit transaction 2> go</pre>6 Disable updates to system tables:

```
1> sp_configure "allow updates", 0
2> go
```

```
1> checkpoint
2> go
```

- 7 Run `dbcc checkcatalog` again to verify that the problem is corrected. If the object entry still exists in other tables, you may get other error messages (such as error 2513). If you do receive other errors, follow the directions in this manual for handling those errors.
- 8 If you wanted the object that has been deleted, re-create it using the information that you obtained in step 3.

Version in which this error is raised

All versions

Error 2520

Severity

16

Message text

Database named %.*s not found; check sysdatabases

Explanation

This error occurs when Adaptive Server is unable to find the database that you specified when you ran `dbcc checkcatalog`, `dbcc checkdb`, or `dbcc dbrepair`.

Possible reasons for this error are:

- Mistyping the database name or the database you specified does not exist
- Corrupt indexes exist on `sysdatabases`

Action

Determine whether the database specified in the error message exists:

```
1> sp_helpdb
2> go
```

Search the output to determine whether you misspelled the database name or whether the `dbid` for the database you specified exists but the name has been removed or changed.

	<p>If the database you are specifying exists and you are getting a 2520 error, run <code>dbcc checkdb</code> on the master database to find out whether it is corrupt. If <code>dbcc checkdb</code> returns error messages, corruption has occurred. Whether or not you believe corruption has occurred, contact Sybase Technical Support. You may be able to rebuild the indexes for <code>master</code> or you may have to restore <code>master</code> from a backup.</p>
Additional information	<p>Before calling Sybase Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the output of <code>dbcc checkdb</code>.</p>
Version in which this error is raised	<p>All versions</p>

Error 2521

Severity	16
Message text	<p>Versions earlier than 15.0:</p> <pre>Table Corrupt: Page is linked but not allocated. Run DBCC TABLEALLOC to correct the problem. (alloc page#=%ld, extent id=%ld, logical page#=%ld, object id in extent=%ld, index id in extent=%ld, object name=%S_OBJID)</pre> <p>Version 15.0 and later:</p> <pre>Table Corrupt: Page is linked but not allocated. Run DBCC %s to correct the problem. (allocation page#=%ld, extent id=%ld, logical page#=%ld, object id in extent=%ld, object name=%S_OBJID, index id in extent=%ld)</pre> <hr/> <p>Note That the first substitution parameter (<code>%s</code>) could be either <code>tablealloc</code> or <code>textalloc</code>. The actual value that displays, <code>tablealloc</code> or <code>textalloc</code>, indicates which <code>dbcc</code> command should be run to correct the problem.</p> <hr/>
Explanation	<p>This error is serious, especially if it occurs on a table’s data pages. Error 2521 indicates that a page is currently in use by a table or index but has not been marked as allocated. If the same page is allocated again, a loss of whatever data resides on the page can result.</p>

Pages that encounter the 2521 error are not included in a database dump. This is because database dumps are performed by reading allocation pages and not by traversing page chains. Therefore, this error should be corrected before dumping the database.

Note The instructions below are for fixing 2521 errors. Two easy-to-use strategies exist for detecting this error earlier. Refer to “Detecting allocation errors early” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery*.

Action

Occasionally `dbcc checkalloc` reports this error when no real error condition exists. Check to see if the error is real, or continue reading this section and take action to correct the condition even if you are unsure whether a real allocation error occurred. This error can be caused by a hardware problem.

Because the process used to determine whether the error is real can be time-consuming, you may want to go directly to “Error resolution” on page 291.

Verifying that the error is real

Run `dbcc checkalloc` in single-user mode if you suspect the 2521 error message are incorrect. If the error is in `master`, use the procedure “Starting Adaptive Server in single-user mode” in the “Server Recovery Tasks” chapter in the most recent version of the guide *Troubleshooting and Disaster Recovery*. Refer to “`dbcc`” in the *Reference Manual: Commands* for information about using `dbcc checkalloc`.

Error resolution

If the error persists, it is possible to clear them all by using `dbcc checkalloc` and `dbcc checkalloc` with the `fix` option. Refer to “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery* for information about using `dbcc checkalloc`.

If the text of the error message includes a real object name, not a number, then the error is on an existing object to which the system catalog has correct references. Continue to “Identify table: User or system table” on page 292.

If a number appears instead of the object name, then that object only partially exists and the error must be corrected using the procedure described in “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of the guide *Troubleshooting and Disaster Recovery*.

Identify table: User or system table

Look at the value for “object ID in extent” in the error message. If the value is 100 or greater, go to “Action for user tables.” If the “object ID in extent” is less than 100, the error occurred on a system table and requires a different procedure described in the section “Action for system tables” on page 293.

Action for user tables

If the “object id in extent” in the error message is 100 or greater, follow these steps to correct the error:

- 1 Check the value of the “index ID in extent” in the error message to determine whether it is a table (the value is 0 (zero)) or an index (with a value between 0 (zero) and 255).
- 2 Run `dbcc tablealloc` or `dbcc indexalloc`, depending on whether the object named in the 2521 error is a table or an index. Before you run the appropriate command, keep in mind:
 - `dbcc tablealloc` can correct this problem on a table or an index, but if the problem is on an index, you can avoid affecting the entire table by using the `dbcc indexalloc` command. If the table is large or heavily used, it may be more practical to use `dbcc indexalloc`.
 - These commands correct the error only when run in the full or optimized mode. Do not specify the `nofix` option or the 2521 error will not be corrected.
 - Use either the “object name” or “object ID in extent” values from the error message in the commands above where the argument “object_name” appears.

Use the command appropriate for your situation:

Tables (index id in extent = 0)	Indexes (0 < index ID in extent < 255)
<pre>1> dbcc tablealloc (<object_name>) 2> go</pre>	<pre>1> dbcc indexalloc (<object_name>, 2> <index_id_in_extent>) 3> go</pre>

Refer to “dbcc” in the *Reference Manual: Commands* and “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about the `dbcc tablealloc` and `dbcc indexalloc` commands.

Action for system tables

If the “object ID in extent” in the error message is less than 100, perform these steps:

1 Put the affected database in single-user mode:

- If the database is *master*, use the procedure “Starting Adaptive Server in single-user mode” in the “Server Recovery Tasks” chapter in the most recent version of the guide *Troubleshooting and Disaster Recovery*, then continue with this procedure.
- If the database is not *master*, use the `sp_dboption` stored procedure to put the affected database in single-user mode:

```
1> use master
2> go
```

```
1> sp_dboption <database_name>, single, true
2> go
```

```
1> use <database_name>
2> go
```

```
1> checkpoint
2> go
```

2 Check the value of the “index ID in extent” in the error message to determine whether it is a table (the value is 0 (zero)) or an index (the value is greater than 0 (zero)).

3 Run `dbcc tablealloc` or `dbcc indexalloc`, depending on whether the object named in the 2521 error message is a table or an index. Before you run the appropriate command, keep in mind:

- `dbcc tablealloc` can correct either a table or an index, but if the problem is on an index, you can avoid affecting the entire table by using the `dbcc indexalloc` command. If you need to minimize the amount of time the table is unavailable, it may be more practical to use `dbcc indexalloc`.
- These commands correct the error only when run in the full or optimized mode, with the `fix` option specified. This is necessary because the default value is `nofix` when these commands are run on system tables.
- Use either the “object name in extent” or “object id” values from the error message in the commands above where the argument “object_name” appears.

Use the command appropriate for your situation:

Tables (index ID in extent = 0)	Indexes (0 < index ID in extent < 255)
<pre>1> dbcc tablealloc (<object_name>, 2> full, fix) 3> go</pre>	<pre>1> dbcc indexalloc (<object_name>, 2> <index_id_in_extent>, full, fix) 3> go</pre>

4 Turn off single-user mode in the database:

- If the database is master, use the procedure in “Returning Adaptive Server to multiuser mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.
- If the database is not master, use this procedure:

```
1> use master
2> go
```

```
1> sp_dboption <database_name>, single, false
2> go
```

```
1> use <database_name>
2> go
```

```
1> checkpoint
2> go
```

Refer to “dbcc” in the *Reference Manual: Commands* and “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about the dbcc tablealloc and dbcc indexalloc commands.

Version in which this error is raised

All versions

Error 2524

Severity

16

Message text

Versions earlier than 15.0:

```
Table Corrupt: Row length is inconsistent between the
computed row length and the recorded row length on page;
check the following page and row: pageno=%ld row#=%d
computed row length=%d row length on page=%ld
```

Version 15.0 and later:

Table Corrupt: Row length is inconsistent between the computed row length (%d) and the row length (%d) recorded on the page. Check row number %d and page number %d, partition ID %d.

Note This error may be caused by a hardware problem.

Explanation Adaptive Server computes the length of each row by using the values stored in the offset and adjust tables of each row. Error 2524 occurs when `dbcc checkdb` or `dbcc checktable` detects that the actual length of a row does not match the computed length of that row.

This error can occur on a data page as well as on an index page.

Action If you have clean backups of your database, recover from backups. If not, refer to “Identifying and fixing a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* to correct the error.

Warning! Some data on this page might be lost if you recover your table using `bcp` or `select into` (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

Before dumping your database, make sure it is working correctly by running the following commands prior to each dump:

- `dbcc checkdb`.
- `dbcc checkalloc` or `dbcc checkalloc` with the `fix` option. (Refer to “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for information about running these commands in multiuser mode and how to prevent spurious allocation errors from `dbcc` commands.)

Hardware errors

Error 2524 is probably caused by a problem within Adaptive Server but may also be caused by hardware errors on your database devices. To help rule out this possibility, do the following:

- Check the Adaptive Server error log to determine whether there are other indications of hardware problems, such as kernel messages reporting I/O errors.

- Check the operating system error log or diagnostics utilities for I/O errors. Refer to “Checking the operating system error log” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.

If problems persist, refer to the error documentation in this manual or contact Sybase Technical Support.

Additional information

Refer to `bcp` in the Adaptive Server *Utility Guide*.

Refer to “select” in the *Reference Manual: Commands* for information about the `select into` command.

Have the information listed in “Reporting errors” on page 4 available before you call Sybase Technical Support, including the output of `dbcc checkdb` and `dbcc checkalloc`.

Version in which this error is raised

All versions

Error 2525

Severity

16

Message text

```
Table Corrupt: Object id wrong; tables: alloc page %ld
extent id=%ld l page#=%ld objid in ext=%ld (name =
%S_OBJID) objid in page=%ld (name = %S_OBJID) objid in
sysindexes=%ld (name = %S_OBJID)
```

Explanation

This error occurs when `dbcc checkalloc` detects one of the following problems:

- The object ID on a page does not match the object ID on the allocation structure (extent) associated with that page. This is the case when the “objid in ext” matches the “objid in page” from the error message.
- The object ID on a page does not match the object ID in the system table `sysindexes`. This is the case when the “objid in ext” does not match the “objid in sysindexes” from the error message.

Action Compare the “objid in ext” to “objid in page” from the error message. If they do not match, go to “Object IDs do not match” on page 297. If they are equal, go to “Object IDs match” on page 298.

Warning! If you drop an object before you correct error 2525, the drop operation could cause pages belonging to another object to be deallocated. This can cause 1108 errors. See “Error 1108” on page 183 for additional information.

Object IDs do not match

There are two possible situations in this scenario. Either the object exists and may be corrupted or it no longer exists. To determine which scenario you have, do the following:

```
1> use <database_name>
2> go

1> select object_name(<objid_in_extent>)
2> go
```

Where *<database_name>* is the database where the corruption is reported and *<objid_in_extent>* is from the error message.

If the result of the query is a table that currently exists or if either object (object ID on the page or the object ID of the allocation structure associated with that page) is a system table (object ID is less than 100), call Sybase Technical Support.

If the query does not return any rows, the object that is marked as owning the extent no longer exists.

If the table is a user table, you may be able to correct this error. Back up the data in the table either use `bcop` to copy out the table whose object ID appears in “objid in page” or by selecting the data into a new table. Refer to “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for instructions.

Once a copy of the table has been made, do one of the following:

- Create a clustered index on the table. If you already have a clustered index, drop it, then re-create it.

or

- Drop the corrupt table, re-create it, and move the data back into the table using a bcp in, an insert with a subselect from the holding table, or select into. Alternatively, rename the holding table once you drop the original table.

Object IDs match

There are two possible scenarios:

- The pointer to a distribution page for a particular index may be invalid.
- There may be a mismatch between the index ID on the allocation page and the index ID on the data page.

Use the following commands to determine which index has the problem. This query compares the “id” column in the output to the “objid in sysindexes” from the error message, and compares the “distribution” column in the output to the “l_page #” from the error message.

```
1> use <database_name>
2> go

1> select indid, name, object_name(id)
2> from sysindexes
3> where id = <objid_in_sysindexes> and
4> distribution = <l_page_#>
5> go
```

If this query does not return any rows, contact Sybase Technical Support.

If the query does return a row, drop and re-create the index specified by this query to correct the problem. (Refer to the *Reference Manual: Commands* for more information about drop index and create index.)

Execute dbcc tablealloc to make sure that the problem is corrected. If the problem still exists, contact Sybase Technical Support.

Additional information

If the object IDs match, dbcc extentdump and dbcc page can also be used to investigate the problem.

Warning! Use the dbcc page command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

Version in which this error is raised All versions

Error 2529

Severity 16

Message text Versions earlier than 15.0:

 Table Corrupt: Attempted to get page %ld, object %ld;
 got page %ld, object %ld.

Version 15.0 and later:

 Table Corrupt: Attempted to get page %d, partition ID
 %d; got page %d, partition ID %d

Note “partition ID” in version 15.0 and later is equivalent to object ID.

Explanation This error occurs when `dbcc checkdb`, `dbcc checktable`, `dbcc checkalloc`, or `dbcc tablealloc` detects an inconsistency between the location of a page on disk and the page number stored in the page header. In the best case, only the value of the page number is wrong. In the worst case, the entire page is corrupt. This error may be caused by a hardware problem.

Warning! This is a serious error that must be corrected immediately.

If the page number in the error is evenly divisible by 256, the error has occurred on an allocation page. All objects referred to by that allocation unit (that page and the 255 pages following it) may be affected (up to 32 objects). Otherwise, the error occurred on a system or user table.

Warning! Attempts to drop an object with a 2529 error may result in error 821 (see “Error 821” on page 127), error 813 (see “Error 813” on page 125), and additional database corruption.

Action 1 Note the object ID displayed in the error message.

 2 Run `dbcc checkalloc` and `dbcc checkdb` on the database affected to determine the full extent of the corruption.

- 3 The 2529 error often results from hardware errors, or system administration problems, such as inadvertent use of a raw partition for two separate purposes. Check on and correct such problems before attempting to restore the database:
 - Refer to “Choosing between raw partitions and UNIX files” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for information about using raw partitions.
 - Check the Adaptive Server error log to determine whether there are other indications of hardware problems, such as kernel messages reporting I/O errors.
 - Check the operating system error log or use the diagnostics utilities for I/O errors.
- 4 If the 2529 error is raised on a system table (the object ID is less than 100), call Sybase Technical Support. If the object ID is 99, another option is to restore from backups.
- 5 If the 2529 error is on a user table (the object ID is greater than or equal to 100), you may be able to select the data into another table and drop the original table. Refer to “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* for instructions.
- 6 After you drop the original table, you may get an 821 error (see “Error 821” on page 127). In this case, shut down Adaptive Server and restart it. Try a select from the old corrupt table to make sure the error is gone, then run `dbcc checkalloc` and `dbcc checkdb` to confirm that there is no more corruption.

In some cases, a 2529 error is first reported as a 2503 error (see “Error 2503” on page 275), and becomes a 2529 error after Adaptive Server is restarted. Although the procedure described above may clear the 2529 error, it will not remove the main cause of your corruption.

Additional information

Before calling Technical Support, have the information listed in “Reporting errors” on page 4 available, including the output of `dbcc checkdb` and `dbcc checkalloc`.

Version in which this error is raised

All versions

Error 2540

Severity

16

Message text

Versions earlier than 15.0:

```
Table Corrupt: Page is allocated but not linked; check
the following pages and ids: allocation pg#=%ld extent
id=%ld logical pg#=%ld object id on extent=%ld (object
name = %S_OBJID) indid on extent=%ld
```

Version 15.0 and later:

Note The following message that displays indicates which dbcc command to run to correct the error.

```
Table Corrupt: Page is allocated but not linked. Run
DBCC %s to correct the problem. (allocation page#=%ld,
extent id=%ld, logical page#=%ld, object id in
extent=%ld, object name=%S_OBJID, index id in
extent=%ld)
```

Explanation

This error occurs when dbcc checkalloc determines that a page is marked as allocated to an object but that page is not being used. There is no corruption or data loss associated with this error.

Each 2540 error means the loss of one blank data page. A few 2540 errors are no cause for concern. However, if many of these errors occur, the amount of “lost” disk space could be significant.

Note The instructions that follow are for fixing 2540 errors once they have occurred. Two easy-to-use strategies exist for detecting this error earlier. Refer to “Fixing and preventing allocation errors” and “Detecting allocation errors early” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for information.

Action

Occasionally dbcc checkalloc reports this error when no real error condition exists. You can either check to determine whether the error is real, or continue with this section and take action to correct it, whether or not it reflects a real allocation error.

Because the process used to discover whether or not the error is real can be time-consuming, you can go directly to “Error resolution” on page 291.

Verifying that the error is real

Run `dbcc checkalloc` in single-user mode if you suspect the 2540 error messages are incorrect. If the error is in `master`, use the instructions in “Starting Adaptive Server in single-user mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* to correct the error. Refer to “`dbcc`” in the *Reference Manual: Commands* for information about `dbcc checkalloc`.

Error resolution

If these errors persist, it is possible to clear them all at once by using the `dbcc checkalloc` and `dbcc checkalloc with fix option` commands. Refer to “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for information about using `dbcc checkalloc`.

If the text of the error message includes a real object name, not a number, then the error is on an existing object to which the system catalog has correct references, and you should continue to “Identify table: User or system table” on page 302.

If a number appears instead of the object name, then that object only partially exists and the error must be corrected using the procedure described in “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery*.

Identify table: User or system table

Look at the value for “`object id on extent`” in the error message. If it is 100 or greater, continue with “Action for user tables.” If the “`object id on extent`” is less than 100, it is a system table and requires the procedure “Action for system tables” on page 303.

Action for user tables

If the “`object id on extent`” in the error message is 100 or greater, follow these steps to correct the error:

- 1 Check the value of the “`indid on extent`” in the error message to determine whether it is a table (value = 0) or an index (value > 0).
- 2 Run `dbcc tablealloc` or `dbcc indexalloc`, depending on whether the object named in the 2540 error message is a table or an index. Before you run the appropriate command, keep in mind:

- dbcc tablealloc corrects this problem on a table or an index, but if the problem is on an index, use dbcc indexalloc to avoid affecting the entire table. If the table is large or heavily used, it may be more practical to use dbcc indexalloc.
- These commands can correct the error only when run in full or optimized mode and with the nofix option not specified (the default for user tables).
- Use the object name or object ID in the following commands where the argument *<object_name>* appears.

Use the command appropriate for your situation:

Tables (index ID in extent = 0)	Indexes (0 < index ID in extent < 255)
1> dbcc tablealloc (<object_name>) 2> go	1> dbcc indexalloc (<object_name>, 2> <indid_on_extent>) 3> go

Refer to “dbcc” in the *Reference Manual: Commands* and “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about dbcc tablealloc and dbcc indexalloc.

Action for system tables

If the “object id on extent” in the error message is less than 100, follow these steps to correct the error:

- 1 Put the affected database in single-user mode:
 - If the database is master, use the procedure in “Starting Adaptive Server in single-user mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*, then go to step 2.
 - If the database is not master, use the sp_dboption stored procedure to put the affected database in single-user mode:

```
1> use master
2> go
```

```
1> sp_dboption <database_name>, single, true
2> go
```

```
1> use <database_name>
2> go
```

```
1> checkpoint
```

```
2> go
```

- 2 Check the value of the “`indid on extent`” in the error message to determine whether it is a table (value is 0) or an index (value is greater than 0).
- 3 Run `dbcc tablealloc` or `dbcc indexalloc`, depending on whether the object named in the 2540 error message is a table or an index. Then execute the appropriate command. Before you run the appropriate command, keep the following in mind:
 - `dbcc tablealloc` corrects either a table or an index, but if the problem is on an index, you can avoid affecting the entire table by using `dbcc indexalloc`. To minimize the amount of time the table is unavailable, use `dbcc indexalloc`.
 - These commands correct the error only when run in the `full` or `optimized` mode, with the `fix` option specified, because the default value is `nofix` on system tables.
 - Use either the “`object name in extent`” or “`object id`” values from the error message in the commands where the argument `<object_name>` appears.

Use the command appropriate for your situation:

Tables (index ID in extent = 0)	Indexes (0 < index ID in extent < 255)
<pre>1> dbcc tablealloc (object_name, 2> full, fix) 3> go</pre>	<pre>1> dbcc indexalloc (object_name, indid_on_extent, 2> full, fix) 3> go</pre>

- 4 Turn off single-user mode in the database:
 - If the database is `master`, refer to “Returning Adaptive Server to multiuser mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.
 - If the database is not `master`, use this procedure:

```
1> use master
2> go
```

```
1> sp_dboption <database_name>, single, false
2> go
```

```
1> use <database_name>
2> go
```

```
1> checkpoint
2> go
```

Refer to “dbcc” in the *Reference Manual: Commands* and “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about dbcc tablealloc and dbcc indexalloc.

Version in which this error is raised

All versions

Error 2546

Severity

16

Message text

Table Corrupt: Extent id %ld on allocation pg# %ld has objid %ld and used bit on, but reference bit off.

Explanation

This error is reported by dbcc checkalloc if an allocation structure or extent is not linked to the other extents for the object referenced by “objid”, but does reference a page that is linked in the object’s page chain. Each occurrence of this error can represent a loss of up to eight pages on disk, or 16K. The pages cannot be used until the error is corrected. Error 2546 can lead to data corruption and can result in various run-time failures. If no other errors are occurring, you can wait until nonpeak hours to correct the problem.

Note The instructions below are for fixing 2546 errors once they have occurred. Two easy-to-use strategies exist for detecting this error earlier. Refer to “Fixing and preventing allocation errors” and “Detecting allocation errors early” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for information about these strategies.

Action

Occasionally dbcc checkalloc reports this error when no real error condition exists. You can either check to determine whether the error is real, or continue with this section and take action to correct it, whether or not it reflects a real allocation error.

Because the process used to discover whether or not the error is real can be time-consuming, you may want to go directly to “Error resolution” on page 306.

Verifying that the error is real

Run `dbcc checkalloc` in single-user mode if you suspect the 2546 error messages are incorrect. If the error is in master, use “Starting Adaptive Server in single-user mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for instructions. Refer to “dbcc” in the *Reference Manual: Commands* for information about `dbcc checkalloc`.

Error resolution

If many of these errors are occurring, it is possible to clear them all at once by using `dbcc checkalloc` and `dbcc checkalloc` with the `fix` option. Refer to “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for information about using `dbcc checkalloc`.

Execute the following query to make sure the object exists and is correctly referred to in the system catalog:

```
1> use <database_name>
2> go

1> select object_name (<objid_from_error_msg>)
2> go
```

If an object name is returned, then the error is on an existing, correctly referenced object. In this case, go to the section “Identify table: User or system table.”

If a number, or something other than an object name is returned, use “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.

Identify table: User or system table

Look at the value for “objid” in the error message. If it is 100 or greater, continue to the next section “Action for user tables.” If the “objid” is below 100, it is a system table and requires a different procedure as described in the section “Action for system tables” on page 307.

Action for user tables

If the “objid” in the error message is 100 or greater, follow these steps to correct the error:

- 1 Run the dbcc page command described in “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for instructions to obtain the value for the “indid.” Substitute the value for “extent id” in the 2546 error message in place of the “page_number” as described in the procedure.
- 2 Run dbcc tablealloc or dbcc indexalloc, depending on whether the value for “indid” indicates it is a table or an index. Before you run the appropriate command, keep in mind:
 - dbcc tablealloc corrects this problem on a table or an index, but if the problem is on an index, use dbcc indexalloc to avoid affecting the entire table. If the table is large or heavily used, it may be more practical to use dbcc indexalloc.
 - These commands will correct the error only when run in full or optimized mode as long as the nofix option is not specified (the default for user tables).
 - Use the object name, if you know it, or the “objid” value from the error message in the commands above where the argument <objid> appears.

Use the command appropriate for your situation:

Tables (index ID in extent = 0)	Indexes (0 < index ID in extent < 255)
1> dbcc tablealloc (<objid>) 2> go	1> dbcc indexalloc (<objid>, <indid>) 2> go

Refer to “dbcc” in the *Reference Manual: Commands* and “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about dbcc tablealloc and dbcc indexalloc.

Action for system tables

If the “objid” in the error message is less than 100, follow these steps to correct the error:

- 1 Run the `dbcc page` command described in the procedure “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* to obtain the value for the “`indid`.” Substitute the value for “`extent id`” in the 2546 error message in place of the “`page_number`” as described in the procedure.

Warning! The procedure in step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 Put the affected database in single-user mode:
 - If the database is `master`, use the procedure in “Starting Adaptive Server in single-user mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*, then go to step 3.
 - If the database is not `master`, use the `sp_dboption` stored procedure to put the affected database in single-user mode:

```
1> use master
2> go
```

```
1> sp_dboption <database_name>, single, true
2> go
```

```
1> use <database_name>
2> go
```

```
1> checkpoint
2> go
```

- 3 Run `dbcc tablealloc` or `dbcc indexalloc`, depending on whether the value of the “`indid`” from step 1 above indicates that it is a table or an index, then execute the appropriate command. Before you run the appropriate command, keep these facts in mind:

- dbcc tablealloc correct either the error on a table or an index, but if the problem is on an index, you can avoid affecting the entire table by using dbcc indexalloc. If you need to minimize the amount of time the table is unavailable, it may be more practical to use dbcc indexalloc.
- These commands will correct the error only when run in the full or optimized mode with the fix option specified, because the default for system tables is nofix.
- You can use the object name if you know it, or “object id” value from the error message in the commands above where the argument <objid> appears.

Use the command appropriate for your situation:

Tables (indid = 0)	Indexes (0 < indid < 255)
1> dbcc tablealloc (<objid>, 2> full, fix) 3> go	1> dbcc indexalloc (<objid>,<indid>, 2> full, fix) 3> go

4 Turn off single-user mode in the database:

- If the database is master, use “Returning Adaptive Server to multiuser mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.
- If the database is not master, use the following procedure:

```
1> use master
2> go
```

```
1> sp_dboption <database_name>, single, false
2> go
```

```
1> use <database_name>
2> go
```

```
1> checkpoint
2> go
```

Refer to “dbcc” in the *Reference Manual: Commands* and “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about dbcc tablealloc and dbcc indexalloc.

Version in which this error is raised

All versions

Error 2547

Severity 16

Message text Table Corrupt: Segment number %d specified for use in the segment map in sysusages has no entry in syssegments. The entry in sysusages is for dbid %d (db name = %S_DBID), with a virtual starting address of %ld.

Explanation Three system tables store information about segments: master..sysusages, and two system tables in the user database, sysindexes and syssegments. When you allocate a device to a database, an entry is added to master..sysusages, and sysusages.segmap provides bitmaps to the segments of the database on each device. When you create a database, Adaptive Server creates the syssegments table in that database and adds rows in the table for the system, default, and log segment. Whenever you add or update segments in a user database, syssegments is updated to reflect new or changed information about the user segments.

Error 2547 is raised when a dbcc check finds an invalid value of sysusages.segmap for a segment, which can happen if you load a user database with user segments different from those in the target database. For example, if you load a database and the target database is larger than the source of the load, and contains a segment that is not in the source, the syssegments table in the newly loaded database will not contain the segment from the higher portion of the target database. As a result, the segmap column in sysusages includes a value not found in syssegments in the newly loaded database, and the dbcc check raises the 2547 error.

Note If the segmaps disagree within the size of the database being loaded, sysusages will be updated to match the loaded database and no 2547 error will be encountered.

Action You can correct the problem by updating sysusages to remove the affected segment information, then defining the necessary segment with sp_addsegment.

Take the following steps:

1 Execute the following command:

```
1> select * from master..sysusages where dbid = <dbid>
2> go
```

Where <dbid> is the dbid in the error message text.

Note that the segment number reported in the message text is the exponent (power) to which 2 is raised to derive `sysusages.segmap`. Use this information to find the `segmap` row in the query output corresponding to the value reported in the error.

For example, if `sysusages` contained a row with `segmap = 11`, and your error text reported segment number 3; then $2^{**}3 = 8$, which is the number by which you need to reduce the `segmap` value (11) to arrive at a corrected `segmap` value of 3.

2 Allow updates to system tables:

```
1> sp_configure "allow updates", 1
2> go
```

3 Correct the affected `sysusages` row:

```
1> update master..sysusages set segmap = <new_segmap>
2> where dbid = <dbid> and segmap = <old_segmap>
3> go
```

In our example:

```
1> update master..sysusages set segmap=3 where dbid=5 and segmap=11
2> go
```

4 Check that exactly one row was affected, then commit:

```
1> commit
2> go
```

5 Disallow updates to system tables:

```
1> sp_configure "allow updates", 0
2> go
```

6 Shut down and restart the server.

7 Run `sp_addsegment` to add any necessary segments to the database. Remember to dump the database after using `sp_addsegment`.

Additional information

If errors persist, contact Sybase Technical Support. Before you call, have the information on hand listed in “Reporting errors” on page 4, including the output:

- `select @@version`
- `select * from master..sysusages where dbid = <db_id>`

Where `<db_id>` is the database ID in the message text.

- `select * from <database_name>..syssegments`

Where *<database_name>* is the database named in the message text.

Segment errors often result from loading a dump to a target database that is logically different (that is, the placement of segments on the devices is different) from the source database. For a detailed discussion of segments and segment mapping, see

- “Creating and Using Segments” in the *System Administration Guide: Volume 2*.
- TechNote 1324, “Segment Remapping with load database When Moving a Database.” This document is available in the Sybase Product Manuals Web site at <http://www.sybase.com/detail?id=1324>.

Version in which this error is raised

All versions

Error 2550

Severity

16

Message text

Missing segment in sysusages segmap.

Explanation

This error is reported by `dbcc checkcatalog` if there is a row in `sysusages` (in the master database) that has a value of 0 (zero) in the `segmap` column. A value of 0 (zero) means that the corresponding space will not be used for further space allocation (data, log, or user-defined objects). The error is not fatal but it does indicate that any free space on that section of the database will not be used.

Possible causes of this error are:

- All of the segments have been removed from a logical device (using `sp_dropsegment`).
- A direct update has been made to the `segmap` column of `sysusages` where the value has been changed to 0 (zero).

Action

If you do not want any new space allocated on the affected segment, no action is necessary.

Otherwise, run `sp_addsegment` or `sp_extendsegment` on the Sybase logical device that has the problem. If the segment does not already exist in `syssegments`, use `sp_addsegment`. If the segment exists, use `sp_extendsegment`.

To add a new segment:

```
1> use <database_name>
2> go

1> sp_addsegment <segname>, <database_name>, <devname>
2> go
```

Where:

- *<segname>* is the name of the new segment.
- *<database_name>* is the name of the database where the segment is to be defined.
- *<devname>* is the name of the database device where *<segname>* will be located. (A database device may have more than one segment associated with it. The space on the new segment will be used only if you create tables or indexes on that new segment.)

To extend an existing segment on the device that has a segmap of 0 (zero):

```
1> use <database_name>
2> go

1> sp_extendsegment <segname>, <database_name>, <devname>
2> go
```

Where:

- *<segname>* is the name of the existing segment. You can extend the segment called default. If you do, the command would be:

```
1> use <database_name>
2> go

1> sp_extendsegment "default", <database_name>,
2> <devname>
3> go
```
- *<database_name>* is the name of the database where the segment is to be extended.
- *<devname>* is the name of the database device that has the segmap 0.

To check that the problem has been resolved, type:

```
1> select * from sysusages where
2> dbid = db_id("<database_name>") and
3> segmap = 0
4> go
```

	Where <i><database_name></i> is the name of the database where the segment you added or extended resides. If no rows are returned, the problem has been resolved.
Additional information	For further information about <code>sp_addsegment</code> and <code>sp_extendsegment</code> , refer to the <i>Reference Manual: Procedures</i> . For more information about segments, refer to “Creating and Using Segments” in the <i>System Administration Guide: Volume 2</i> .
Version in which this error is raised	All versions

Error 2558

Severity	16
Message text	<p>Versions earlier than 15.0:</p> <pre>Extent not within segment: Object %ld, indid %d includes extents on allocation page %ld which is not in segment %d.</pre> <p>Version 15.0 and later:</p> <pre>Extent not within segment: Object %ld, indid %d, partition %ld includes extents on allocation page %ld which is not in segment %d</pre>
Explanation	<p>A segment is a label used to point to one or more database devices. Segments map the future allocations of different types of data to different devices.</p> <p>Error 2558 occurs when <code>dbcc checkalloc</code>, <code>dbcc tablealloc</code>, or <code>dbcc indexalloc</code> discovers that a portion of a database resides on a segment which was not defined to contain that type of data. For example, if <code>dbcc checkalloc</code> discovers user data on a segment designated to hold transaction log data, error 2558 will occur.</p>

Note Error 2558 is only raised when the 2513 trace flag is turned on.

When a database gets this error, it still functions. However, 1105 errors may occur as a side effect of the 2558 error.

In addition, if you have log on your data segment, and you lose the data device for some reason, you cannot get a good transaction dump with `no_truncate`. This means you cannot get up-to-the-minute recovery.

Sybase recommends that you do not turn on trace flag 2513 unless you have been getting 1105 errors or have recently run `sp_logdevice`, `sp_placeobject`, or alter database commands. When trace flag 2513 is on, `dbcc checkalloc`, `dbcc tablealloc`, and `dbcc indexalloc` commands take significantly longer to run and a lot of additional output may be created.

To turn on trace flag 2513, use these commands:

Pre-15.0.2:

```
1> dbcc traceon (2513)
2> go
```

Version 15.0.2 and later:

```
1> set switch on 2513 with override
2> go
```

To turn off trace flag 2513, use these commands:

Pre-15.0.2:

```
1> dbcc traceoff (2513)
2> go
```

Version 15.0.2 and later

```
1> set switch off 2513
2> go
```

Creating segments

By default, when Adaptive Server creates a database, it creates the following three segments:

- **system segment** – allows system catalog tables to be placed on devices that contain this segment for a particular database.
- **default segment** – allows user tables or indexes to be placed on devices that contain this segment without specifying the `on` clause (for example, `create table test (i int) on segment_1`) for a particular database.

- log segment – allows the transaction log to be placed on devices that contain this segment for a particular database.

Note A device can contain more than one type of segment, although a segment can contain only one type of data (user data, transaction log data, and so on). Thus, different data can be on the same device, but not on the same segment.

For example, this command creates the default, system and log segments for the database `example`:

```
1> create database example
2> on device_1 = 10
3> log on device_2 = 2
4> go
```

The system and default segments are mapped to `device_1` and the log segment is mapped to `device_2`. For the `example` database, all future space allocations for system or user objects will occur on `device_1` and all transaction log records will be placed on `device_2`. Other databases may use other portions of `device_1` or `device_2` and these devices may have different segment mappings for that database.

A 2558 error would result in the `example` database if system or user objects were located on `device_2` or transaction log records on `device_1`.

Causes of error 2558

Following is a summary of how and when error 2558 is raised. For details, refer to the sections listed in the “See...” column in the following table.

Causes of error 2558	See...
Loading into a database that has different data and log mapping than the dumped database.	“Different data and log mapping”
When you use <code>sp_placeobject</code> , the old allocation area is still in use ^a .	“ <code>sp_placeobject</code> ” in <i>Reference Manual: Procedures</i>
When you use <code>sp_dropsegment</code> , existing objects are still mapped to the dropped segment.	“ <code>sp_dropsegment</code> ” in <i>Reference Manual: Procedures</i>
a. Error 2558 is only raised in this case when trace flag 2513 is turned “on.”	

Different data and log mapping

If this error occurred after you re-created and loaded a database from a dump, there were probably incorrect or different entries in the `master.sysusages` table. A database loaded from a dump must be created in exactly the same way as the database that was dumped. If it is not, different entries in `master.sysusages` can cause 2558 errors.

This can cause a “data on log” situation, where data, such as user or system objects, is loaded into a log segment. This means that less total log space is available to the database. The data is still accessible, but processing may be slowed, if not totally stopped, by the shortage of log space.

There can also be a “log on data” situation, where portions of the transaction log are loaded into a data segment. This is not as serious as the “data on log” scenario. This is because the portion of the transaction log on the data segment is eventually truncated using the `dump transaction` command, freeing up the space. All future transaction log allocations will be correctly placed on the log segment.

Suppose the `example` database were altered as follows:

```
1> alter database example on device_1 = 2
2> go
```

The logical order of the database pages is:

- The first 10MB of pages are data (default and system segments)
- The next 2MB are log
- The last 2MB are data

A 2558 error could result if the database were dumped, then dropped and re-created with different segment mappings. For example:

```
1> create database example on device_1 = 12
2> log on device_2 = 2
3> go
```

Although both ways of creating the database (create and alter, or just create) allocate the same amount of data and log space on the same devices, the mapping of data and log space is not the same. The first 12MB are allocated differently: in the first (create and alter) database, 10MB of data are followed by 2MB of log; in the second (create), the first 12MB are data. Dumping the first database and loading it to the second would cause 2MB of log to be mapped onto a data segment, yielding 2558 errors.

Warning! If you load databases in this manner, you can get mapping that you do not expect and that can cause problems. Do not load databases in this manner! If you suspect that such a problem might have occurred, use `sp_helpdb` to check the mapping of your database.

All space allocations that occur on Database 2 after the load completes are correctly mapped.

sp_placeobject

If you use `sp_placeobject` to assign an object to a new segment, existing pages will still reside on the old segment.

If you have trace flag 2513 turned “on,” `dbcc checkalloc`, `dbcc tablealloc`, and `dbcc indexalloc` will display error 2558 after you use `sp_placeobject`.

sp_dropsegment

If you have trace flag 2513 turned “on,” `dbcc checkalloc`, `dbcc tablealloc`, and `dbcc indexalloc` will display error 2558 if you use `sp_dropsegment` to remove a segment for which existing objects are still mapped.

Prevention

- Consider using the `on segment_name` option with `create table` rather than using `sp_placeobject` to assign an object to a segment.
- Do not drop segments when objects are still mapped to them.
- Keep an up-to-date copy of the `master..sysusages` system table as well as scripts to re-create all databases. Use the scripts to ensure that the entries in the `master..sysusages` table for the database being loaded match the corresponding entries for the database that was dumped. More specifically, the `segmap`, `lstart`, and `size` columns of `sysusages` must be identical in content and order. Verify this before beginning the load database command.
- The entries in `sysusages` will be correct on the database if you execute the following commands with the same parameters and in the same order as they were on the dumped database:

- create database
- alter database
- sp_addsegment
- sp_dropsegment
- sp_extendsegment
- sp_logdevice

Note You cannot depend on the output of the system procedure `sp_helpdb` to compare the contents of the `sysusages` table, because it does not necessarily display rows from `sysusages` in the actual logical page order. Instead, do direct selects against the `sysusages` table.

Action

You only need to resolve 2558 errors if the object taking up space on the wrong segment creates a problem for you.

If you received the error because of an `sp_dropsegment`, you may be able to resolve it using `sp_addsegment` or `sp_extendsegment`.

If there are a substantial number of 2558 errors due to an incorrect load database, you can drop and correctly re-create the database, then reload it.

Use this section if you have 2558 errors and cannot re-create the database from scripts or hard copy.

2558 errors occur on these types of pages:

- Data pages or clustered index pages
- Nonclustered index pages
- text or image data pages
- System tables
- Transaction log pages

The action needed to correct this error depends on the type of page on which the error occurred.

Use the following table to match up the object ID value (“Object” in the error message) and index ID (“indid” from the error message) with the corresponding data page type, then go to that section for the appropriate action. All action sections follow the table.

Type of data	Corresponding object ID and index ID values
“Data pages or clustered index pages”	Object ID > 99, index ID = 0 or 1

Type of data	Corresponding object ID and index ID values
“Nonclustered index pages”	Object ID > 99, 1 < index ID < 255
“text or image data pages”	Object ID > 99, index ID = 255
“System table pages”	Object ID < 100 (and not equal to 8)
“Transaction log pages”	Object ID = 8

An object ID value of 100 will not occur.

Data pages or clustered index pages

(Object ID > 99 and Index ID = 0 or 1)

The error occurred on the data page or on the clustered index of a user table. Resolve the problem with either one of the following procedures:

- Create a clustered index on the table, if you have the available space needed. If a clustered index already exists, drop and re-create it. If not, creating one will clear the error. You can then drop the clustered index.
- Copy the data out of the table, drop and re-create the table, and copy the data back in using `bcp` or `select into`.

Nonclustered index pages

(Object ID > 99 and 1 < Index ID < 255)

The error occurred on the nonclustered index of a user table. The error can be cleared by dropping and re-creating the index. Creating a clustered index causes all nonclustered indexes to be rebuilt as well. Therefore, if 2558 errors are occurring on several indexes on the same table, you can create a clustered index on that table to clear all these errors.

text or image data pages

(Object ID > 99 and index ID = 255)

The error occurred on data which is text or image datatype. Use either procedure below to correct the situation:

- 1 Select all the information from the old table into a new table.
- 2 Drop the old table. Any subsequent `dbcc checkalloc`, `dbcc tablealloc`, or `dbcc indexalloc` command will continue to show this error until the old table is dropped.
- 3 Use `sp_rename` to rename the new table with the old table name.

or:

Copy the table out, using `bcp` or `select/into`. Then drop and re-create the table, and copy the contents of the table back in, using `bcp` or `select/into`.

System table pages

(Object ID < 100 and not 8)

Call Sybase Technical Support.

Transaction log pages

(Object ID = 8)

If the error occurs on pages containing transaction logs, do the following:

- 1 Have all users finish transactions and wait until this procedure is finished to initiate any further transactions.
- 2 Perform the checkpoint command:

```
1> use <database_name>
2> go

1> checkpoint
2> go
```
- 3 Dump the transaction log according to your normal procedures.
- 4 Notify users that they may resume normal operations.

Additional information

Refer to “Creating and Using Segments” in the *System Administration Guide: Volume 2* for information about segments and devices.

Refer to the *Reference Manual: Commands* for information about checkpoint and dump transaction.

Version in which this error is raised

All versions

Error 2559

Severity

16

Message text

```
Data page number %ld is empty but is not the first page.
Status = 0x%x.
```

Explanation When an object is first created, a single, empty page is allocated to it. Usually there are no other empty pages. Error 2559 occurs when `dbcc checktable` or `dbcc checkdb` encounters an empty page while traversing the page chain of an object which is not the first page of the chain.

If you are not encountering any other errors either at run time or from `dbcc`, it is possible that the empty page will not cause any further problems. Therefore, it is acceptable to wait to correct this problem until nonpeak hours. If other errors are occurring, refer to the information in this manual if applicable or contact Sybase Technical Support.

Action Correct this situation by forcing the page chain involved to be rebuilt. This can be accomplished in different ways, depending on whether or not the page chain involved is associated with an index or table data. To determine this, refer to the procedure “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* to determine the index ID and object ID associated with the page number from the error message.

Warning! “Finding an object name from a page number” uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

If the error is on a user table (a system table’s object ID is less than 100), go to “User tables or clustered indexes (index ID is 0 or 1)” on page 322.

If the error is on a system table and the index ID is not 0 (zero), the error is on a system table index. Refer to “Fixing a corrupted system table index” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* for instructions on how to repair the system table index.

If the error is on a system table and the index ID is 0 (zero), contact Sybase Technical Support. They may be able to help you repair the corruption or it may be necessary to restore from clean backups.

User tables or clustered indexes (index ID is 0 or 1)

There are two ways to correct this error, depending on whether you have enough extra disk space to create a clustered index on the table in question.

Solution 1: Drop index

If a clustered index exists, drop and re-create it. If not, create a dummy clustered index and then drop it. This will remove the problem since creation of a clustered index causes the table to be copied and does not copy the empty pages. Note that clustered index creation does require considerable disk space (refer to “create index” in the *Reference Manual: Commands*).

Solution 2: Use the bulk copy (bcp) utility

Use bcp to copy the table out of and then back to the Server:

- 1 bcp the table data out to an operating system file.
- 2 Empty the table with the truncate table command.
- 3 (Optional) Drop all indexes on the table.
- 4 bcp the table data into Adaptive Server.
- 5 (Optional) Re-create all indexes on the table.

For nonclustered indexes (1 < index ID < 255)

There are two ways to correct this error on a nonclustered index:

Solution 1: Drop index

- 1 Execute the following query in the database in question to determine the index name:

```
1> select name from sysindexes
2> where id = <object_ID> and indid = <index_ID>
3> go
```

Where <object_ID> and <index_ID> are from the output of dbcc page that you ran earlier.

- 2 Drop and re-create the nonclustered index specified in the query.

Solution 2: Create a Clustered Index

Create a clustered index since this forces all nonclustered indexes to be rebuilt. This method may not be suitable for large tables because of the overhead required for creating a clustered index.

Version in which this error is raised

All versions

Error 2572

Severity	10
Message text	Database '%S_DBID' is not in single user mode - may find spurious allocation problems due to transactions in progress.
Explanation	<p>This warning occurs when you execute <code>dbcc checkalloc</code> on a database that is not in single-user mode. <code>dbcc checkalloc</code> displays this message and then continues to execute.</p> <p>Executing <code>dbcc checkalloc</code> on a database that is not in single-user mode may cause error messages to display that do not reflect the actual state of your database. These error messages are referred to as spurious (or false). Error messages 2521, 2540, 2546, 7939, 7940, and 7949 may be spurious under these conditions.</p>
Action	<p>No action is required.</p> <p>If other errors are reported by <code>dbcc checkalloc</code>, refer to this manual for information on how to deal with them.</p> <p>If your operational environment allows it, put the questionable database into single-user mode prior to running <code>dbcc checkalloc</code>.</p> <ul style="list-style-type: none">• For the master database, refer to “Starting Adaptive Server in single-user mode” in the chapter “Server Recovery Tasks” in the most recent version of <i>Troubleshooting and Disaster Recovery</i>.• For other databases, use <code>sp_dboption</code> as follows:<pre>1> use master 2> go 1> sp_dboption <database_name>, single, true 2> go 1> use <database_name> 2> go 1> checkpoint 2> go</pre>
Additional information	Refer to “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of <i>Troubleshooting and Disaster Recovery</i> .

Refer to “dbcc” in the *Reference Manual: Commands* for information about dbcc checkalloc.

Version in which this error is raised

All versions

Error 2573

Severity

16

Message text

Database '%.*s' is not marked suspect. You cannot drop it with DBCC.

Explanation

This error occurs when you try to use dbcc dbrepair to drop a database that is not marked suspect. Only databases that have been marked suspect can be dropped using this command.

Databases are marked suspect in the following ways:

- Adaptive Server marks a database suspect when some critical errors occur.
- The Sybase System Administrator can mark a database suspect in order to drop it with the dbcc dbrepair command.

Action

Use drop database to try to drop the database. If this fails, use the procedure described in “Dropping a database when drop database fails” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery*.

Version in which this error is raised

All versions

Error 2574

Severity

16

Message text

Versions earlier than 15.0:

Index page number %ld is empty. Status = 0x%x.

Version 15.0 and later:

Index page number %d, partition ID %d, is empty. Status = 0x%x

Explanation	<p>Tables can contain text columns, which are variable-length columns that can hold printable characters, and image columns, which are variable-length columns that can hold binary data. Adaptive Server stores text and image data in a linked list of data pages separate from the rest of the table. Each 2K text or image page stores a maximum of 1800 bytes of data.</p> <p>Error 2574 occurs when a dbcc check finds problems with a page that is part of a table's text or image page chain. Error 7105 is a related error; it is raised when Adaptive Server attempts data retrieval or another operation on a table with a bad text or image page chain.</p>
Action	<p>Error 2574 is typically due to a small number of bad rows in a table with text or image data, and can be corrected by dropping and re-creating those rows.</p> <p>Use the following steps to address the problem:</p> <ol style="list-style-type: none">1 Identify the table. Refer to “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of <i>Troubleshooting and Disaster Recovery</i> for instructions about how to use the dbcc page command to determine the table. <hr/> <p>Warning! The procedure referenced in step 1 uses the dbcc page command. Use the dbcc page command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.</p> <hr/> <ol style="list-style-type: none">2 Determine the first page of the text or image chain using dbcc pglinkage, using as the starting page the page number from the error message. For example, assuming that dbid=7 and the starting page number is 531: <pre>1> dbcc pglinkage(7, 531, 0, 1, 0, 0) 2> go Object ID for pages in this chain = 208003772. Beginning of chain reached. Page : 531 Page : 532 Page : 530 3 pages scanned. Object ID = 208003772. Last page in scan = 530.</pre> <p>Page 530 is the first page of the text or image chain.</p>

- 3 Identify the rows with the bad data using the following command:

```
1> select <column1> from <table_name>
2> where convert (int, textptr(<text_colname>))
3> = <first_page_of_chain>
4> go
```

Where *<column1>* is any column or columns in your table that will uniquely identify the rows.

- 4 Drop the rows in question.
- 5 Insert the rows back into the table.
- 6 Run `dbcc checktable` on the table to verify that the problem has been corrected.

Version in which this error is raised

All versions

Error 2575

Severity

16

Message text

Versions earlier than 15.0:

The last page %ld in sysindexes for table '%.*s' has next page # %ld in its page header. The next page # should be NULL. Please check sysindexes.

Version 15.0 and later:

The last page of table '%.*s' is recorded in sysindexes as page number %d, partition ID %d. In its header, this page has page number %d as its 'next page'; however, the 'next page' value should be NULL. Check sysindexes.

Explanation

The `sysindexes` table lists each table and index, and the segment where each table, clustered index, nonclustered index, and chain of text pages is stored. It also lists other information such as the first page and root for the object. For heap tables (those with no clustered index), the root entry points to the last page in the table's data page chain.

Adaptive Server expects to find a value of zero for the next page number on the last page for a table (meaning that there are no more pages in the page chain). Error 2575 is raised when `dbcc checktable` or `dbcc checkdb` is run, and the next page number in the table's last page is not zero. This error is caused by an incorrect root page value for the table in `sysindexes`, or corruption in the table's last page.

Action

If this is the only error message you see when you run `dbcc checktable` or `dbcc checkdb`, it is likely that the table's root entry in `sysindexes` is incorrect. Use one of these options to recover from the problem:

- Select the table data into a new table, drop the old table, and rename the new table to the old table name. Run `dbcc checktable` to check that the problem has been resolved.
- Bulk copy the affected table out, drop and re-create the table, and bulk copy back in. This is the most efficient solution for a very large table. Run `dbcc checktable` to confirm that the problem is resolved.

For more information about how to copy a table in a new table or file, refer to “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.

- Call Sybase Technical Support. They may be able to determine the correct last page number and patch the table's root entry in `sysindexes`.

If you see other errors in addition to error 2575, the problem may be more extensive. Follow the instructions in this manual for handling those errors.

Additional information

Before calling Sybase Technical Support, have the information listed in “Reporting errors” on page 4 available, including the output of `dbcc checktable` or `dbcc checkdb`.

Version in which this error is raised

All versions

Error 2578

Severity

16

Message text

The first page %ld in `sysindexes` for table '%.*s' has previous page # %ld in its page header. The previous

page # should be NULL. Please check *sysindexes*.

Note In Adaptive Server Enterprise version 15.0 and later, error 9923 (“The first page %ld in *syspartitions* for partition %ld of table '%.*s' has previous page # %ld in its page header. The previous page # should be NULL”) has replaced error 2578.

Explanation

The *syspartitions* table lists each table and index partition and the segment where the partition is stored. *syspartitions* maintains key storage management information, such as the first page of a data or index page chain, the last page of a heap, the root page of an index partition, and so on.

Adaptive Server expects to find a value of zero for the previous page number on the first page for a table (meaning that there are no prior pages in the page chain). Error 2578 is raised when *dbcc checktable* or *dbcc checkdb* is run, and the previous page number in the table’s first page is not zero. This error is caused by an incorrect first page value for the table in *sysindexes*, or corruption in the table’s first page.

Action

If the object with the error is a system table, call Sybase Technical Support. They may be able to help you repair the corruption, but it may be necessary to restore from clean backups.

If the table is a user table, and this is the only error message you see when you run *dbcc checktable* or *dbcc checkdb*, it is likely that the table’s first page entry in *sysindexes* is incorrect. Use one of these options to recover from the problem:

- Select the table data into a new table, drop the old table, and rename the new table to the old table name. Run *dbcc checktable* to check that the problem has been resolved.
- Bulk copy the affected table out, drop and re-create the table, and bulk copy back in. This is the most efficient solution for a very large table. Run *dbcc checktable* to confirm that the problem is resolved.

For more information about how to copy a table in a new table or file, refer to “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.

If you see other errors in addition to error 2578, the problem may be more extensive. Follow the instructions in this manual for handling those errors.

Additional information	Before calling Sybase Technical Support, have the information listed in “Reporting errors” on page 4 available, including the output of dbcc checktable or dbcc checkdb.
Version in which this error is raised	Version 12.5.x and earlier, and error 9923 in version 15.0 and later

Error 2582

Severity 16

Message text

Versions earlier than 15.0:

Data page (%S_PAGE) has been marked as an overflow page; however, the previous page (%S_PAGE) does not indicate that there is an overflow page linked to it.

Version 15.0 and later:

Data page %S_PAGE, partition ID %d, has been marked as an overflow page; however, the previous page (page %S_PAGE, partition ID %d) does not indicate that an overflow page is linked to it.

Explanation

When using a clustered index, Adaptive Server maintains data so that the, physical order of rows is the same as their logical, indexed order. The bottom, or leaf level, of a clustered index contains the table’s actual data pages. When an index key has a large number of duplicates, all of the data rows may not fit on the same data page. Due to the structure of a clustered index, it is not possible to simply insert these overflow rows to the beginning of the next data page. Instead, Adaptive Server allocates and inserts a new data page between the original data page and the next data page, marking this an overflow page to store the additional values (and marking the original page as having an overflow page). The original page, and the overflow page linked to it, now both contain at least one row with that key value. The overflow page contains only rows with an overflow index value.

On tables without a clustered index, the second page through the last page of the table are marked as overflow pages; the first page through the next-to-last page are marked as having an overflow page.

When dbcc checktable examines a page marked as an overflow page, it checks the previous (original) page for consistency. Error 2582 is raised when the original page is not marked as having an overflow page.

Action	<p>Take the following steps to correct the problem:</p> <ol style="list-style-type: none"> 1 Use the procedure described in “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of <i>Troubleshooting and Disaster Recovery</i> to translate the second page number displayed in the error message into an object ID. <hr/> <p>Warning! The procedure referenced in step 1 uses the <code>dbcc page</code> command. Use the <code>dbcc page</code> command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.</p> <hr/> <ol style="list-style-type: none"> 2 If there is a clustered index on the table, run <code>dbcc tablealloc</code> to correct the problem: <pre style="margin-left: 40px;">1> dbcc tablealloc (<table_name>, full, fix) 2> go</pre> <hr/> <p>Note The <code>tablealloc</code> command corrects the error only when you run it in full or optimized mode.</p> <hr/> <ol style="list-style-type: none"> 3 If there is no clustered index on the table, use one of these options to recover from the problem: <ul style="list-style-type: none"> • Select the table data into a new table, drop the old table, and rename the new table to the old table name. Run <code>dbcc checktable</code> to confirm that the problem has been resolved. <li style="text-align: center; padding: 5px 0 5px 20px;">or • Bulk copy the affected table out, drop and re-create the table, and bulk copy back in. This is the most efficient solution for a very large table. Run <code>dbcc checktable</code> to confirm that the problem is resolved.
Additional information	For more information about how to copy a table in a new table or file, refer to “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> .
Version in which this error is raised	All versions

Error 2583

Severity	16
Message text	<p>Versions earlier than 15.0:</p> <p>Number of data pages counted by DBCC differs from count in the data OAM.</p> <p>Version 15.0 and later:</p> <p>Number of data pages (%ld) counted by DBCC differs from count in the data OAM (%ld) for dbid %d object %ld.</p>
Explanation	<p>The object allocation map (OAM) tracks information about the allocation and deallocation of storage space for objects (tables and indexes). An OAM entry exists for each table (except syslogs and sysgams) and for the table's indexes.</p> <p>All pages for an object are logically chained together, each page containing a header that includes the page numbers of the previous and next pages in the chain.</p> <p>Error 2583 is raised when Adaptive Server detects that the count of OAM data is different from the actual number of pages in the chain. This problem is detected by <code>dbcc checkdb</code> and <code>dbcc checktable</code>.</p> <p>Usually, this error will not cause operations to fail at run time. Therefore, it is acceptable to wait to correct the problem until nonpeak hours.</p> <p>Error 2583 is similar to error 7949; both are raised due to OAM data inconsistencies. Error 2583 occurs when Adaptive Server finds incorrect used page counts for the object, whereas error 7949 occurs when the server just finds incorrect unused page counts.</p> <hr/> <p>Note The instructions below are for fixing 2583 errors once they have occurred. Two easy-to-use strategies exist for detecting this error earlier. Refer to “Detecting allocation errors early” in the chapter “Other Useful Tasks” in the most recent version of <i>Troubleshooting and Disaster Recovery</i>.</p> <hr/>
Action	<p>If many of these errors are raised, as might be the case if you run <code>dbcc checkdb</code> and several tables are affected, it is possible to clear them all by using the <code>dbcc checkalloc</code> command with the <code>fix</code> option. Refer to “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide for information about using <code>dbcc checkalloc</code>.</p>

If only a small number of these errors are raised and you know the table name (that is, if you ran `dbcc checktable`), determine whether the affected table is a user table or a system table:

```
1> select id from sysobjects
2> where name = <object_name>
3> go
```

User tables have object IDs of 100 or more; system tables have IDs under 100. Corrective action is different in the two cases because the database must be in single-user mode before correcting system tables.

Action for user tables

If the affected table is a user table, run `dbcc tablealloc` to correct the error:

```
1> dbcc tablealloc (<table_name>, full, fix)
2> go
```

Keep in mind that `dbcc tablealloc` can correct the error only when run in the full or optimized mode.

Refer to “`dbcc`” in the *Reference Manual: Commands* and “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about `dbcc tablealloc`.

Action for system tables

If the affected table is a system table, follow these steps to correct the error:

- 1 Put the affected database in single-user mode:
 - If the database is `master`, refer to “Starting Adaptive Server in single-user mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*, then go to step 2.
 - If the database is not `master`, use the `sp_dboption` stored procedure to put the affected database in single-user mode:

```
1> use master
2> go
```

```
1> sp_dboption <database_name>, single, true
2> go
```

```
1> use <database_name>
2> go
```

```
1> checkpoint
```

```
2> go
```

2 Run `dbcc tablealloc` to correct the error:

```
1> dbcc tablealloc (<table_name>, full, fix)
2> go
```

Keep in mind that for a system table, `dbcc tablealloc` can correct the error only when you run it in the full or optimized mode, and with the `fix` option specified.

3 Turn off single-user mode in the database:

- If the database is `master`, refer to “Returning Adaptive Server to multiuser mode” in the in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.
- If the database is not `master`, use this procedure:

```
1> use master
2> go
```

```
1> sp_dboption <database_name>, single, false
2> go
```

```
1> use <database_name>
2> go
```

```
1> checkpoint
2> go
```

Refer to “`dbcc`” in the *Reference Manual: Commands* and “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about `dbcc tablealloc`.

Version in which this error is raised

All versions

Error 2591

Severity

16

Message text

```
Could not find index row which has id %d for table
'%.*s'.
```

Explanation Error 2591 is raised when you run a `dbcc` command if Adaptive Server is unable to find an index in `sysindexes`. The `'id %d'` parameter in the message text shows the index ID of the index that the `dbcc` command expected to find.

Error 2591 occurs with the following `dbcc` commands:

- `dbcc reindex`
- `dbcc indexalloc`

Error 2591 is raised with the following states:

State	Meaning
1	<code>dbcc reindex</code> is unable to locate a primary index row for a clustered index on a system table in <code>sysindexes</code> .
2	<code>dbcc reindex</code> is unable to locate a primary index row for a non-clustered index on a system table in <code>sysindexes</code> .
3	<code>dbcc indexalloc</code> is unable to locate the specified index in <code>sysindexes</code> .

Action If the error was raised when using `dbcc indexalloc`, check the index name supplied in the command, and retry using the correct index name. To verify the available indexes for the table, type:

```
1> sp_help <table_name>
2> go
```

If the error was raised when using `dbcc reindex`, use an alternative method for correcting the index. Note the object name and index id appearing in the message text, and identify the object id as follows:

```
1> select object_id('<table_name>')
2> go
```

Then use the procedure described in “Fixing a corrupted system table index” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for instructions on how to repair the system table index.

Version in which this error is raised

All versions

Error 2596

Severity 16

Message text Versions earlier than 15.0:

```
%S_PAGE has an incorrect pgfreeoff value of %d. The
offset should be %d.
```

Version 15.0 and later:

```
Page %S_PAGE, partition ID %d, has an incorrect
'pgfreeoff' value of %d. The offset should be %d.
```

Explanation

The free byte offset on a page indicates the offset of the first unused byte on that page. Error 2596 occurs when dbcc checktable detects that the free byte offset on a page is incorrect.

Warning! This is a serious error that indicates possible corruption on the affected page and can occur in conjunction with other errors, particularly error 2506.

Action

If error 2596 occurs on a system table being processed by dbcc checktable, call Sybase Technical Support.

Case 1 – Free byte offset is lower

If the number of the free byte offset shown in the error message is lower than the expected number, some data has been lost. To determine how much data has been lost, find the difference (in bytes) between the actual free byte offset and the expected free byte offset (both numbers are shown in the error message output). The amount of data lost includes 2 bytes overhead per row where there are only fixed-length columns or $4 + \langle n \rangle$ bytes overhead per row with variable-length columns, where $\langle n \rangle$ is the number of variable-length columns. If you need help, call Sybase Technical Support.

Decide whether you should restore the database that contains the affected table from backups. In making that decision, consider the number of rows lost and the number of other corruption errors that have occurred on this table and database.

Case 2 – Free byte offset is higher

If the free byte offset on the page header is higher than the expected value, no data has been lost. However, some extraneous rows may exist.

If you have clean backups of your database, recover from backups. If you do not have clean backups of your database, follow these steps to correct the error:

- 1 Use the procedure “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* to identify the table and index correspond to the page number from the error message text.

Warning! The procedure referenced in step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 If the object with the error is not a system table (a system table’s object ID is less than 100), continue with step 3.

If the object with the error is a system table and the index ID is not 0, refer to “Fixing a corrupted system table index” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for instructions on how to repair the system table index.

If the object with the error is a system table and the index ID is 0 (zero), contact Sybase Technical Support. They may be able to help you repair the corruption or it may be necessary to restore from clean backups.

- 3 If the object with the the error is a user table and if the index ID is 0 (zero) or 255, go to step 4.

If the index ID is not 0 (zero) or 255:

- a Translate it into an index name:

```
1> use <database_name>
2> go

1> select name from sysindexes
2> where id = <object_ID> and indid = <index_ID>
3> go
```

- b To ensure that the information needed to re-create the index is available, run `sp_helpindex` on the index prior to dropping it.
- c Drop the index.
- d Re-create the index. This clears the corruption in most cases.

- e Run dbcc checktable on the table to confirm that the corruption is gone.
- 4 If the index ID is 255, delete or update the data row to remove the reference to the bad text page.

If the index ID is 0 (zero), do one of the following:

- Restore the database from clean backups.
- or
- Refer to “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.

Warning! Some data on the page can be lost if you recover the table using bcp or select into (that is, the corrupted row and rows following it can be truncated and contain the wrong keys). Compare the two tables (old and new) row-by-row (for example, by joining them on a primary key) to determine if any rows are different (corrupted).

Before dumping the database, confirm that it is working correctly by running the following commands prior to each dump:

- dbcc checkdb.
- dbcc checkalloc or dbcc checkalloc with the fix option. (Refer to “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for information about running these commands in multi-user mode and how to prevent false allocation errors from dbcc commands.)

Version in which this error is raised

All versions

Error 7901

Severity

16

Message text

Versions earlier than 15.0:

Page %ld was expected to be the first page of a TEXT/IMAGE value.

Version 15.0 and later:

Page %d, partition ID %d, is expected to be the first page of a text/image/unitext value.

Explanation

Adaptive Server stores text and image data in a linked list of pages separate from the rest of the table. Error 7901 is raised when a dbcc utility is processing a text or image page chain, but finds that the first page in the chain is not a text or image page. The error is due to a corrupt page header in the first page, or a bad text pointer in a table row.

Error 7901 is raised with these states:

State	Meaning
1	dbcc checktable detected the first page error while checking a table's text or image chains.
2	dbcc fix_text detected the first page error while recalculating the statistics for a table's text pages.
3	dbcc rebuild_text detected the first page error while rebuilding the internal data structures for text or image data. (This state is raised in version 12.0 and later).

Action

Take these esteps to resolve the error.

1 Identify the rows with the bad data using the following command:

```
1> select <unique_col> from <table_name>
2> where convert (int, textptr(<text_colname>))
3> = <text_page>
4> go
```

Where *<unique_col>* is any column or columns in the table that will uniquely identify the rows, *<table_name>* is the name of the table, *<text_colname>* is the text column, and *<text_page>* is the page ID in the message text.

2 Drop the rows in question.

3 Run dbcc checktable on the table to confirm that the problem has been corrected.

Alternately, you can create a new table using `select into`, specifying all rows from the original table except the affected rows identified in step 1. “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for details.

Version in which this error is raised

All versions

Error 7902

Severity 16

Message text Versions earlier than 15.0:

Data size mismatch occurred while checking TEXT/IMAGE values. The first page for this value is: %ld. There were %ld bytes found, which is different from the expected data length of %ld bytes.

Version 15.0 and later:

Data size mismatch while checking text/image/unitext values. The first page for this value is %d, partition ID %d. %d bytes found; however, the expected data length is %d bytes.

Note This error may be caused by a hardware problem.

Explanation This error occurs when Adaptive Server detects a mismatch between the expected length of the text or image data of a text column and its actual length. The length of the text or image data is stored on the page header of the text or image data page.

This error usually occurs in one of the following situations:

- When you run `dbcc checktable` on a table whose text or image column was updated to NULL and then updated again to a non-NULL value via a `writetext` command which did not commit.
- When you run `dbcc checktable` or `dbcc checkdb` and your text or image data is corrupted. If this is the case, check your hardware error log and your operating system error log and determine the cause of the 7902 error.

Prevention Consider using `writetext with log` or `update` instead of `writetext with no log`. This way, you will not have unlogged changes to your database and will not get 7902 errors if `writetext with no log` did not commit.

Action To eliminate the 7902 error, perform the following steps for each table encountering it:

- 1 Record the page numbers reported by error 7902 for your table.
- 2 Select the offending rows from the table in a separate temporary table, giving the name of the text column for the `textptr` function:

```
1> select * into #<newtable> from <mytable>
2> where convert (int, textptr(<text_colname>))
```



```
3> in (<list_of_page#'s>)
4> go
```

The list of page numbers in the above query is the enumeration of all the page numbers specified by the 7902 errors. *<text_colname>* is the name of the text column.

- 3 Update the corrupted text or image data from the old table with the text or image data from the new table:

```
1> begin transaction
2> go
1> update <mytable>
2> set <mytable>.<text_col> = #<newtable>.<text_col>
3> from #<newtable>
4> where <mytable>.<other_col> =
5>#<newtable>.<other_col>
6> go
```

Verify that only the required number of rows are updated by the above query, then commit the transaction. *<other_col>* in the above query is the name of another column in the table that uniquely identifies a row (like a primary key). If you do not have any primary keys on the old table, run the above query with several search conditions in the where clause:

```
1> update <mytable>
2> set <mytable>.<text_col> = #<newtable>.<text_col>
3> from #<newtable>
4> where <mytable>.<col1> = #<newtable>.<col1>
5> and <mytable>.<col2> = #<newtable>.<col2>
6> go
```

- 4 Make sure that only the required number of rows are updated by the above query and then commit the transaction:

```
1> commit transaction
2> go
```

Example

You will see the following messages during dbcc checktable:

```
Checking mytable
Msg 7902, Level 16, State 1:
Server 'PRODUCTION', Line 1:
Data size mismatch occurred while checking
TEXT/IMAGE values. The first page for this value
is: 321. There were 12 bytes found, which is
different from the expected data length of 0 bytes.
```

Msg 7902, Level 16, State 1:

Server 'PRODUCTION', Line 1:

Data size mismatch occurred while checking TEXT/IMAGE values. The first page for this value is: **291**. There were 17491 bytes found, which is different from the expected data length of 19291 bytes.

Run the following query to select the erroneous rows from `<mytable>` into `#<newtable>`:

```
1> select * into #<newtable> from <mytable>
2> where convert (int, textptr(<text_colname>))
3> in (321, 291)
4> go
```

Alternate method to eliminate error 7902

An alternate method is to select all data from the old table into a new table.

If the index ID is 0 or 255 and you do not have clean backups of your database, copy the data from the corrupted table into a new (dummy) table or into a file. Then rename your old, corrupted table and copy the data back into a new table using the original name. For more information about doing this, refer to “Rescuing data from a corrupted table” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for details.

Warning! Some data on this page might be lost if you recover your table using `bcp` or `select into` (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

Before dumping your database, make sure it is working correctly. More specifically, the following commands should be run prior to each dump:

- `dbcc checkdb`.
- `dbcc checkalloc` or `dbcc checkalloc` with the `fix` option. Refer to “`dbcc`” in the *Reference Manual: Commands* for information about `dbcc checkalloc`.

Additional information

For more information, refer to “`select into`”, “`convert`”, and “`writetext`” in the *Reference Manual: Commands* and in the *Transact-SQL User's Guide*.

Version in which this error is raised

All versions

Error 7928

Severity	16
Message text	Index %.*s is not consistent; found %ld leaf rows but %ld data rows. Drop and recreate the index.
Explanation	This error is returned by <code>dbcc checktable</code> or <code>dbcc checkdb</code> (which calls <code>dbcc checktable</code>). When checking the integrity of a nonclustered index, <code>checktable</code> compares the leaf row count (the total number of rows in the leaf pages for the index) with the count of data rows. Error 7928 is raised when the leaf row count does not match the data row count, and indicates an index problem.
Action	<p>Correct the problem by re-creating the index as follows:</p> <ol style="list-style-type: none">1 Check the output of the <code>dbcc</code> query to identify the table associated with the index.2 If the table is a system table, refer to “Fixing a corrupted system table index” in the chapter “Other Useful Tasks” in the most recent version of <i>Troubleshooting and Disaster Recovery</i> for instructions on how to repair the system table index, then go to step 4.3 If the table is a user table, take these steps:<ul style="list-style-type: none">• Run <code>sp_helpindex</code> on the index to ensure that the information needed to re-create the index is available.• Drop the index.• Re-create the index.4 Run <code>dbcc checktable</code> on the table to verify that the corruption is gone. If the corruption still exists, call Sybase Technical Support.
Version in which this error is raised	All versions

Error 7930

Severity	16
Message text	<p>Versions earlier than 15.0:</p> <p>Table Corrupt: keys in left child do not precede the parent key; check left child page %ld.</p>

Version 15.0 and later:

Table Corrupt: Keys in the left child page do not precede the parent key; check page %d, partition ID %d.

Explanation

This serious error is returned by `dbcc checktable` or `dbcc checkdb` (that calls `dbcc checktable`). `dbcc checktable` checks a table and its indexes to determine whether:

- Pages are linked correctly
- Indexes are properly sorted
- Pointers are consistent
- Page offsets are consistent
- `sysindexes` pages are correct
- The data row count is equal to the leaf row count for nonclustered indexes

During index-level checking, for each key in the parent page, `dbcc checktable` checks whether the last key in the left child page is less than the parent key, and whether the first key in the right child page is greater than or equal to the parent key. If the last key in the left child is not less than the parent key, error 7930 occurs.

Action

- 1 Use the procedure in “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* to identify which table and index correspond to the page number from the error message text.

Warning! The procedure referenced in step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 If the object encountering the error is not a system table (a system table’s object ID is less than 100), go to step 3.

If the object with the error is a system table and the index ID is not 0 (zero), refer to “Fixing a corrupted system table index” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for instructions on how to repair the system table index. then go to step 5.

If the object with the error is a system table and the index ID is 0 (zero), contact Sybase Technical Support. They may be able to help you repair the corruption but it may be necessary to restore from clean backups.

- 3 If the object with the error is a user table and the index ID is not 0, determine whether there is a clustered index on the table:

```
1> sp_helpindex <table_name>
2> go
```

Where *<table_name>* is the name you determined in step 1.

If there is a clustered index on the table, rebuild it, then go to step 5.

- 4 If there is a nonclustered index on the table, rebuild the index using these steps.

- a Translate the index ID into an index name:

```
1> use <database_name>
2> go

1> select name from sysindexes
2> where id = <object_ID> and indid = <index_ID>
3> go
```

- b Drop the index.

- c To ensure that the information needed to re-create the index is available, run `sp_helpindex` on the index prior to dropping it.

- d Re-create the index. This clears the corruption in most cases.

- 5 Run `dbcc checktable` on the table to verify that the corruption is gone. If the corruption still exists, call Sybase Technical Support.

Additional information

Refer to “drop index” and “create index” in the *Reference Manual: Commands* for information about dropping and re-creating indexes.

Version in which this error is raised

All versions

Error 7939

Severity 16

Message text Versions earlier than 15.0:

Table Corrupt: The entry is missing from the OAM for object id %ld indid %d for allocation page %ld.

Version 15.0 and later:

Table Corrupt: The entry is missing from the OAM for object id %ld indid %d partition %ld for allocation page %ld.

Explanation

This error is raised when Adaptive Server detects that the allocation pages associated with an object are not recorded in the object allocation map (OAM) for the object. This problem is detected by `dbcc checkalloc`, `dbcc tablealloc`, and `dbcc indexalloc`.

Usually, error 7939 will not cause any operations to fail at run time, so it is acceptable to wait to correct this problem until non-peak hours. However, do not drop a table with the 7939 error, as subsequent `dbcc checkalloc` commands may produce 2540 and 2546 errors.

Note The instructions below are for fixing 7939 errors once they have occurred. Two easy-to-use strategies exist for detecting this error sooner in the future. Refer to “Detecting allocation errors early” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for information about these strategies.

Action

Occasionally `dbcc checkalloc` reports this error when no real error condition exists. You can either check to see if these errors are real, or continue with this section and take action to correct them, whether or not they reflect a real allocation error.

Because the process used to discover whether or not the errors are real can be time-consuming, you may want to go directly to the “Error resolution” section now.

Verifying that the error is real

Run `dbcc checkalloc` in single-user mode if you suspect the 7939 error messages are incorrect. If the error is in `master`, use the section “Starting Adaptive Server in single-user mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* to start Adaptive Server in single-user mode. Refer to “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about `dbcc checkalloc`.

Error resolution

If many of these errors are occurring, it is possible to clear them all at once by using the `dbcc checkalloc` with `fix` option command. Refer to “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for information about using `dbcc checkalloc`.

Confirming that the object exists

Execute the following query to make sure the object exists and is correctly referred to in the system catalog:

```
1> use <database_name>
2> go

1> select object_name (<objid_from_error_msg>)
2> go
```

If an object name is returned, then the error is on an existing, correctly referenced object. If this is the case, use the procedure in “Fixing allocation errors when object ID and index ID are known” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery*.

If a number, or something other than an object name is returned, use the procedure in “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery*.

Version in which this error is raised

All versions

Error 7940

Severity 22

Message text The counts in the OAM are incorrect. This implies that there are entries missing. Run tablealloc utility with the FIX option on the table with the inaccurate OAM counts.

Explanation This error is raised when Adaptive Server detects that the total page count in the object allocation map (OAM) is different from the actual number of pages in the chain. This problem is detected by `dbcc checkalloc`, `dbcc tablealloc`, and `dbcc indexalloc`.

This error can result in various run-time failures and should be corrected; however, if no other errors are occurring, you can wait to correct the problem until non-peak hours.

Note The instructions below are for fixing 7940 errors once they have occurred. Two easy-to-use strategies exist for detecting this error sooner in the future. Refer to “Detecting allocation errors early” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for information about these strategies.

Action Occasionally `dbcc checkalloc` reports this error when no real error condition exists. You can either check to see if these errors are real, or continue with this section and take action to correct them, whether or not they reflect a real allocation error.

Because the process used to discover whether or not the errors are real can be time-consuming, you may want to go directly to the “Error Resolution” section now.

Verifying that the error is real

Run `dbcc checkalloc` in single-user mode if you suspect the 7940 error messages are incorrect. If the error is in `master`, use the section “Starting Adaptive Server in single-user mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for instructions about how to invoke Adaptive Server in single-user mode. Refer to “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about `dbcc checkalloc`.

Error resolution

If many of these errors are occurring, it is possible to clear them all at once by using the `dbcc checkalloc` with `fix` option command. Refer to “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* for information about using `dbcc checkalloc`.

Identify table: User or system table

Get the table name that the error occurred on from the original output which indicated this error. If it is a user table, continue with “Action for user tables.” If it is a system table, it requires a different procedure described in “Action for system tables.”

Action for user tables

Run `dbcc tablealloc`. This command can correct the error only when run in the full or optimized mode, and with the `nofix` option not specified (the default for user tables):

```
1> dbcc tablealloc (<table_name>)
2> go
```

Refer to “dbcc” in the *Reference Manual: Commands* and “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about `dbcc tablealloc`.

Action for system tables

1 Put the affected database in single-user mode:

- If the database is master, use the procedure in “Starting Adaptive Server in single-user mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*, then go to step 2.
- If the database is not master, use the `sp_dboption` stored procedure to put the affected database in single-user mode:

```
1> use master
2> go
```

```
1> sp_dboption <database_name>, single, true
2> go
```

```
1> use <database_name>
2> go
```

```
1> checkpoint
2> go
```

- 2 Run `dbcc tablealloc`. This command corrects the error only when run in the full or optimized mode, with the `fix` option specified, because the default value is `nofix` on system tables:

```
1> dbcc tablealloc (<table_name>, full, fix)
2> go
```

- 3 Turn off single-user mode in the database:

- If the database is master, use “Returning Adaptive Server to multiuser mode” in the chapter “Server Recovery Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery*.
- If the database is not master, use the following procedure:

```
1> use master
2> go
```

```
1> sp_dboption <database_name>, single, false
2> go
```

```
1> use <database_name>
2> go
```

```
1> checkpoint
2> go
```

Refer to “`dbcc`” in the *Reference Manual: Commands* and “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about `dbcc tablealloc`.

Version in which this error is raised

All versions

Error 7949

Severity

16

Message text

Versions earlier than 15.0:

```
The number of pages used and unused for %ld index %d on
allocation page %ld do not match the counts in the OAM
entry.
```

Version 15.0 and later:

```
The number of pages used and unused for object %ld index
%d partition %ld on allocation page %ld do not match the
counts in the OAM entry. Actual used/unused counts: [%d,
%d]. Used/unused counts in OAM entry: [%d, %d].
```

Explanation

This error occurs when Adaptive Server detects that the allocation page count in the object allocation map (OAM) is different from the actual number of pages in the chain. This problem is detected by `dbcc checkalloc`, `dbcc tablealloc`, and `dbcc indexalloc`. The object ID is the value immediately before the index value in the error text.

Usually this error will not cause any operations to fail at run time. Therefore, it is acceptable to wait to correct this problem until non-peak hours.

Note The instructions below are for fixing 7949 errors once they have occurred. Two easy-to-use strategies exist for detecting this error sooner in the future. Refer to “Detecting allocation errors early” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for information about these strategies.

Action

Occasionally `dbcc checkalloc` reports this error when no real error condition exists. You can either check to see if these errors are real, or continue with this section and take action to correct them, whether or not they reflect a real allocation error.

Because the process used to discover whether or not the errors are real can be time-consuming, you may want to go directly to the “Error Resolution” section now.

Verifying that the error is real

Run `dbcc checkalloc` in single-user mode if you suspect the 7949 error messages are incorrect. If the error is in master, use the section “Starting Adaptive Server in single-user mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for instructions about how to invoke Adaptive Server in single-user mode. Refer to “Checking Database Consistency” in the *System Administration Guide: Volume 2* for information about `dbcc checkalloc`.

Error resolution

If many of these errors are occurring, it is possible to clear them all at once by using the `dbcc checkalloc` command with the `fix` option. Refer to “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* for information about using `dbcc checkalloc`.

Confirming that the object exists

Execute the following query to make sure the object exists and is correctly referred to in the system catalog:

```
1> use <database_name>
2> go

1> select object_name (<objid_from_error_msg>)
2> go
```

If an object name is returned, then the error is on an existing, correctly referenced object. If this is the case, use the procedure in “Fixing allocation errors when object ID and index ID are known” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery*.

If a number, or something other than an object name is returned, use the procedure in “Fixing and preventing allocation errors” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery*.

Version in which this error is raised

All versions

Error 7989

Severity

16

Message text

The serial allocation flag was found to be improperly set in allocation unit %ld, at extent ID %ld, with alloc map = %d and objid = %ld. This flag has been cleared by DBCC.

Explanation	<p>When you initialize a database device, the space is divided into allocation units of 256 2K pages each. When an object requires space, Adaptive Server allocates a block of 8 2K pages within the unit, known as an extent, to the object.</p> <p>The first page of each allocation unit is an allocation page, which stores information about each extent contained in the allocation unit. The allocation page also contains a <i>serial allocation flag</i> that indicates whether extents in the unit are allocated in serial, unbroken order. This flag enables more efficient processing: dumps and loads, for example, can process a set of transaction log pages as a group without having to traverse individual log pages.</p> <p>dbcc checkalloc examines each allocation page for consistency. When it finds an allocation page with serial allocation turned on, checkalloc verifies that all extents in that allocation unit are indeed allocated serially. If this is not the case, checkalloc raises error 7989.</p>
Action	Error 7989 is informational. No action is needed. checkalloc resets the flag to indicate that the extents are not serially allocated.
Version in which this error is raised	All versions

Error 9961

Severity	17
Message text	DBCC failed to connect to the requested number of worker processes. The number of workers configured is insufficient, or the number of workers exceeds 128. Use sp_dbcc_updateconfig to reduce 'max worker processes' for this database.
Explanation	dbcc checkstorage processing may require the use of one or more worker processes. Error 9961 is raised when the server is unable to obtain the worker processes necessary for a checkstorage run.

Error 9961 is raised with the following states:

State	Meaning
1	There is insufficient memory available to configure more worker processes.
2	checkstorage is unable to start up a worker process.
3	checkstorage is unable to attach to a worker process.

State	Meaning
4	Similar to state 3.
5	Similar to state 3.

Action	<p>1 Determine the resources necessary to run dbcc checkstorage:</p> <pre>1> use master 2> go 1> sp_plan_dbccdb <target database> 2> go</pre> <p>Write down the process count recommended for the database.</p> <p>2 Verify that the following parameter settings provide the number of worker processes needed. Worker process configuration is determined by:</p> <ul style="list-style-type: none"> • <code>max worker processes</code> – specifies the maximum number of worker processes that checkstorage uses for each target database; you set this parameter with <code>sp_dbcc_updateconfig</code>. • <code>number of worker processes</code> – defines the total number of worker processes supported by the server; you set this parameter with <code>sp_configure</code>. Note that worker processes are not dedicated to checkstorage operations. <p>3 Retry the checkstorage operation.</p>
Additional information	See the chapter “Checking Database Consistency” in the <i>System Administration Guide: Volume 2</i> for details on configuring worker processes for checkstorage.
Version in which this error is raised	All versions

Error 12964

Severity	20
Message text	DBCC cannot insert the row in the dbcc_operation_log table for this check of database '%!%', aborting check.
Explanation	The <code>dbcc_operation_log</code> table records the use of the dbcc checkstorage and dbcc checkverify operations.

Error 12964 is raised when the current checkstorage or checkverify cannot insert rows corresponding to the operating database (on which checkstorage and checkverify is active) into the table `dbcc_operation_log`. This can occur if the `dbccdb` database does not exist or the table could be in an inconsistent state, possibly due to the content change in the `dbccdb` database with version 12.5.0.3.

Action	If <code>dbccdb</code> does not exist, create it by running the <code>installdbccdb</code> script found in the <code>\$(SYBASE)\$(SYBASE_ASE)/scripts</code> directory.
Version in which this error is raised	Version 12.5.1 and later

Error 15056

Severity	16
Message text	The root page %1! of index ID %2! for empty table '%3! can have at most one entry pointing to the first page %4!. However, it has either an incorrect entry, or more than one entries.
Explanation	In an empty table with an all-pages-locked (APL) clustered index, the root page should normally have only one entry pointing to the first data page. Similarly, in an empty table with an APL non-clustered index, the root page should normally be the same as first index page and have no entries or the root page may have one entry pointing to an index page. In any case, it should not have more than one entry. This error indicates that the root page was found to have more than one entry and is raised when the inconsistency is detected by <code>dbcc checktable</code> or <code>dbcc checkdb</code> .
Action	To correct the problem, drop and recreate the affected index.
Additional information	If the problem persists, contact Sybase Technical Support. If possible, have the details on hand about how this table is used, which may be helpful in determining the cause of the error.
Version in which this error is raised	Version 12.0.0.8 and later

Error 15080

Severity	16
Message text	<p>The root page %ld of index id %d for table '%.*s' has inconsistent rootkeys information. Drop and recreate the index.</p> <p>Use DBCC REINDEX to rebuild indexes on system catalogs with this form of inconsistency.</p>
Explanation	<p>This error occurs when dbcc checktable finds the root page of a btree index is not coherent with the information in sysindexes.</p>
Action	<p>As noted in the message text:</p> <ul style="list-style-type: none">• If the error occurred on a user table, the index needs to be dropped and re-created.• If the error occurred on a system table, fix the index with dbcc reindex.
Version in which this error is raised	Version 12.5.x and later

Error 15082

Severity	18
Message text	<p>In versions earlier than 15.0.2:</p> <p>There are %d OAM entries with zero used/unused page counts. Run DBCC %s FIX on objid %d indid %d.</p> <p>In version 15.0.2 and later:</p> <p>There are %d OAM entries with zero used/unused page counts. This is not an error. Run DBCC %s FIX on objid %d indid %d to remove these entries as doing so will improve OAM scan performance.</p>
Explanation	<p>This is an informative message (not an error) printed while running dbcc checkalloc, dbcc tablealloc, dbcc indexalloc, or dbcc textalloc. If any of these commands find an object allocation map (OAM) page with a zero used or unused count, this message prints for further action to improve the OAM scan performance.</p>

Action	Run <code>dbcc checkalloc</code> , <code>dbcc tablealloc</code> , <code>dbcc indexalloc</code> , or <code>dbcc textalloc</code> with the <code>fix</code> option to remove the unnecessary entries from the OAM page. The <code>fix</code> option reclaims the space in the OAM page to store new OAM entries.
Additional information	The OAM page is used to track all pages allocated to a table or an index. In addition to the allocation information, it also keeps track of the number of used and unused pages for the table or index. For more information about the OAM, see “Understanding the object allocation map (OAM)” in the <i>System Administration Guide: Volume 2</i> .
Version in which this error is raised	Version 12.5.4 and later

Error 15084

Severity	18
Message text	The OAM has no room for new inserts. Run DBCC %s FIX on objid %d indid %d.
Explanation	This is an informative message (not an error) printed while running <code>dbcc checkalloc</code> , <code>dbcc tablealloc</code> , <code>dbcc indexalloc</code> , or <code>dbcc textalloc</code> . This message prints when the task finds an object allocation map (OAM) page that less than the internally defined threshold value for free insert space. Normally the Adaptive Server OAM manager verifies that there is enough space for a new OAM entry. If this fail, a new OAM page is allocated and linked to the existing OAM pages. At the same time, there might be some unnecessary entries (see “Error 15082” on page 356) on the OAM page that can be reclaimed by running the any of the above <code>dbcc</code> commands with <code>fix</code> option.
Action	Run <code>dbcc checkalloc</code> , <code>dbcc tablealloc</code> , <code>dbcc indexalloc</code> , or <code>dbcc textalloc</code> with the <code>fix</code> option to remove unnecessary entries from the OAM page. The <code>fix</code> option reclaims space in the OAM page to store new OAM entries. This is not a severe error, just a warning that conditions may lead to error 1127 on this object. See “Error 1127” on page 186 for more information.
Additional information	For more information about OAM pages, refer to “Understanding Page and Object Allocation Concepts” in the <i>System Administration Guide: Volume 2</i> .
Version in which this error is raised	Version 12.5.4 and later

This section contains error messages for Adaptive Server Enterprise insert commands.

Error	Page
Error 2601	359
Error 2610	360
Error 2615	361
Error 2619	365
Error 2620	366
Error 2626	368
Error 2628	369
Error 2630	371
Error 2631	372

Error 2601

Severity	14
Message text	Attempt to insert duplicate key row in object '%.*s' with unique index '%.*s'%S_EED
Explanation	No two rows can have the same index value (including NULL) in a column or columns with a unique index. Adaptive Server checks for duplicate values when the index is created (if data already exists) and checks each time data is added with an insert or update. Error 2601 occurs when you try to put duplicate index values into a column or columns with a unique index.
Action	Using a unique index makes sense only when uniqueness is a characteristic of the actual data. Choose one of the following solutions, depending on whether you need a unique index: <ul style="list-style-type: none"> • If you need duplicate index values in indexed columns, drop the unique index and create a nonunique index instead.

- If you need a unique index on data that contains duplicate values, you must change some values to remove the duplicates:
 - a Use a select statement to find the row that will be duplicated by the update or insert command.
 - b Modify either the data in the table or the data that you want to insert, so that the index values do not match.

Additional information

Refer to “create index” in the *Reference Manual: Commands* for information.

Version in which this error is raised

All versions

Error 2610

Severity

22

Message text

```
Could not find leaf row in nonclustered index '%.*s'  
page %ld that corresponds to data row from logical data  
page %ld, row offset %d during update index attempt  
after data page split in dbid %d.
```

Note This error can be caused by a hardware problem.

Explanation

This error occurs when Adaptive Server fails to update a nonclustered index at the leaf level after a page split (a page split occurs when a page fills up and there is no more room to write on that page). This error message indicates index corruption.

This is a serious error. You may be able to recover from this error, but you may have to restore from known, clean backups.

Action

- 1 Record the value of the data page and index name (or ID if that is given instead) from the error message. Using this information and the instructions in “Finding an object name from a page number” in the *Troubleshooting and Disaster Recovery* guide, determine which table corresponds to the data page number from the error message text.

- 2 If the table determined in step 1 is a system table (that is, the object ID is less than 100) and the index ID is not 0 (zero), see “Fixing a system table corrupted index” in the *Troubleshooting and Disaster Recovery* guide for instructions on how to repair the system table index. Then go to step 5.

Note If the table determined in step 1 is a system table and the index ID is 0 (zero), contact Sybase Technical Support. They may be able to help you repair the corruption, but it may be necessary to restore from clean backups.

- 3 If the table determined in step 1 is not a system table and the message displayed the index ID rather than the index name, find the name of the index involved by executing this query:

```
1> select name from sysindexes
2> where indid = <index_ID>
3> and id = <object_ID>
4> go
```

Where *<object_ID>* is the ID of the table as determined in step 1 and *<index_ID>* is the ID of the index specified in the error message.

- 4 Drop and re-create the index using the table name obtained in step 1 and the index name determined in step 3.
- 5 To verify that all problems have been resolved on this table, run `dbcc checktable` and `dbcc tablealloc` on the affected table.

Additional information

If problems persist, this procedure may not be sufficient to clean up the index corruption, and you should contact Sybase Technical Support. Have the output from the `dbcc checkallocc` or `dbcc tablealloc` and `dbcc checktable` commands available for review.

Version in which this error is raised

All versions

Error 2615

Severity

14

Message text

Attempt to insert duplicate row in table '%.*s' with index '%.*s' in database '%.*s'. Could drop and recreate index with ignore duprow or allow duprow.

Explanation

Tables with clustered indexes do not allow duplicate rows. Error 2615 occurs:

- When you try to insert a duplicate row in a table that already has a clustered index.
- When you try to upgrade to Adaptive Server Enterprise version 12.0.x or version 12.5.x, and encounter duplicate rows in the `syskeys` system table. This occurs because the upgrade exposes a previously hidden column, “`spare1`”, setting its contents to 0 (zero). If the database previously contained two or more `syskeys` rows that were identical except for this column, attempting to set `syskeys.spare1` to zero will produce 2615 errors, accompanied by 3462 errors (“Database ‘%. *s’: upgrade could not install optional upgrade item ‘%d’. Please refer to previous error messages to determine the problem. Upgrade will continue.”).

Note The duplicate `syskeys` entries are not caused by the upgrade process; rather, the upgrade process reveals such entries if they exist.

Action

Case 1

If the duplicate values are a result of incorrect data, change the data and rerun the transaction. Otherwise, decide whether to allow or prevent duplicate rows in the table where the clustered index will be created. Use one of the following procedures.

Note `ignore_dup_row` and `allow_dup_row` are mutually exclusive options of the `create index` command. Do not use them together in the same statement, which will cause error 1916.

Allowing duplicate rows

Create the index with the `allow_dup_row` option:

```
1> create clustered index <index_name> on <table_name>
2> with allow_dup_row
3> go
```

Subsequent update and insert commands can create duplicate rows after the `allow_dup_row` option is set.

Warning! Do not use this procedure if you have a unique index on any column or columns in the table. Unique indexes are only used on tables where no duplicate rows exist.

Deleting duplicate rows

Create the index with the `ignore_dup_row` option:

```
1> create clustered index <index_name> on <table_name>
2> with ignore_dup_row
3> go
```

Existing duplicate rows will be deleted from the table as the index is created. Subsequent attempts to enter a duplicate row with `insert` or `update` are ignored and the `insert` or `update` is cancelled with an informational message. If the `insert` or `update` is part of a larger transaction, the rest of the transaction will continue as usual.

Case 2

If the error occurs during upgrade, begin by making sure that you are upgrading using one of the recommended migration paths:

- Pre-12.0 to 12.0.x
- Pre-12.0 or 12.0.x to 12.5.0.1 (interim release) or later

If you do not upgrade using one of these paths, error 2615 may be raised during a required upgrade step and that the error will abort the upgrade. By following the recommended path, you may encounter the error, but it is benign and will not cause the upgrade to abort.

After the upgrade completes, use the following steps after the upgrade to correct the problem by removing duplicate `syskeys` rows:

- 1 Identify duplicate rows.

Execute the following query in the database named in the 2615 error to identify the duplicate `syskeys` entries:

```
1> select id,
2> 'Primary object' = object_name(id),
3> type,
4> depid,
5> 'Dependent object' =
```

```
object_name(isnull(depid,0)),
6> spare1
7> from syskeys k1
8> where exists (select spare1
9> from syskeys k2
10> where k1.id = k2.id
11> and k1.type = k2.type
12> and k1.spare1 != k2.spare1)
13> go
```

Identify the duplicate `syskeys` entries that have two identical (id, type) pairs.

To see if there are more than two entries with identical (id, type) pairs, execute this query:

```
1> select * from syskeys
2> where id = <test_id> and type = <test_type>
3> go
```

The duplicates are the rows where ALL columns except `spare1` match.

2 Eliminate duplicate `syskeys` entries.

a Allow updates to system tables:

```
1> sp_configure "allow updates", 1
2> go
```

b Delete all but one of each duplicate set. Repeat the following command once for each duplicate:

```
1> delete syskeys
2> where id = <duplicate_id>
3> and type = <duplicate_type>
4> and spare1 = <one_of_the_duplicates>
5> go
```

c (Optional) Before disallowing updates to system tables (step d), you can execute the following commands to prevent other duplicates:

```
1> update syskeys
2> set spare1 = 0
3> where spare1 != 0
4> go
```

d After removing all duplicates, disallow updates to system tables:

```
1> sp_configure "allow updates", 0
2> go
```

Additional information	Refer to “create index” in the <i>Reference Manual: Commands</i> for information about <code>ignore_dup_row</code> and <code>allow_dup_row</code> .
Version in which this error is raised	All versions

Error 2619

Severity	20
Message text	Keys did not match overflow page when inserting row at end of page (%S_PAGE) that has overflow page (%S_PAGE) that is not disconnected in dbid = %d.
Explanation	<p>With a clustered index, Adaptive Server maintains data so that the physical order of rows is the same as their logical, indexed order. The bottom or leaf level of a clustered index contains the table’s actual data pages. When an index key has a large number of duplicates, not all of these may fit on the same data page. Adaptive Server then allocates a new data page, marking it an “overflow page” to store the additional values. The original page and the overflow page now both contain at least one row with that key value.</p> <p>If the data corresponding to the key value on the first page is now updated or deleted so that the key no longer appears on that page, Adaptive Server removes the association between the two pages. The overflow page is now a regular page just like any other data page.</p> <p>Error 2619 is raised when Adaptive Server attempts to insert a key value at the end of a page, and the key currently in place has an associated overflow page, but the key does not match the key value in the overflow page. It is likely that the page served as an overflow page at some point, but when a row with that key value on the first page was updated or deleted, the overflow page was not marked as a regular (disconnected) page.</p> <p>When error 2619 occurs, the connection to Adaptive Server is broken. This is a serious error that indicates a corrupt clustered index on the table.</p>
Action	<p>Use the following steps to resolve the error:</p> <ol style="list-style-type: none">1 Perform the procedure in “Finding an object name from a page number” in the most recent <i>Troubleshooting and Disaster Recovery</i> guide to identify the table.2 Use <code>sp_helpindex</code> to list all indexes associated with the table and identify the clustered index on the table.

	<ol style="list-style-type: none">Drop and re-create the clustered index.Run <code>dbcc tablealloc</code> on the table to ensure that the problem is resolved. If errors are raised, contact Sybase Technical Support for assistance.
Additional information	When calling Sybase Technical Support, have the information available that is listed in “Reporting errors” on page 4.
Version in which this error is raised	All versions

Error 2620

Severity	21
Message text	The offset of the row number at offset %d does not match the entry in the offset table of the following page: %S_PAGE.
Explanation	An offset table is stored at the end of every data page and it stores where rows are located on that page. Adaptive Server uses the offset table to quickly access the rows in each page. Error 2620 occurs during an <code>insert</code> operation when Adaptive Server detects inconsistencies in the offset table of the page specified by the error message.

Note This error can be caused by a hardware problem.

This is a serious error. You may be able to recover from this error but you may have to restore from known, clean backups.

Action	<ol style="list-style-type: none">Record the value of the page from the error message. Perform the procedure in “Finding an object name from a page number” in the most recent <i>Troubleshooting and Disaster Recovery</i> guide to determine the table and index that correspond to the page number.If the object encountering the error is not a system table (the object ID is less than 100), continue to step 3.
--------	---

If the object is a system table and the index ID is not 0 (zero), use the instructions in “Fixing a system table corrupted index” in the *Troubleshooting and Disaster Recovery* guide to repair the system table index.

Note If the index ID is 0 (zero), contact Sybase Technical Support. They may be able to help repair the corruption, but it may be necessary to restore from clean backups.

- 3 For user tables, if the index ID is 0 (zero) or 255, continue to step 4.

If the index ID is not 0 (zero) or 255:

- a Translate it into an index name:

```
1> use <database_name>
2> go

1> select name from sysindexes
2> where id = <object_ID> and indid = <index_ID>
3> go
```

- b To ensure that the information needed to re-create the index is available, run `sp_helpindex` on the index prior to dropping it.
- c Drop the index.
- d Re-create the index. This clears the corruption in most cases.
- e Run `dbcc checktable` on the table to verify that the corruption is gone.

- 4 If the index ID is 0 (zero) or 255, do one of the following:

- Restore the database from clean backups.
- or
- Perform the procedure in “Rescuing data from a corrupted table” in the most recent version of the *Troubleshooting and Disaster Recovery* guide.

Warning! Some data on the page can be lost if you recover the table using `bcp` or `select into` (that is, the corrupted row and rows following it might be truncated and contain the wrong keys). Compare the two tables (old and new) row by row (by joining them on a primary key, for example) to determine which rows are different (corrupted).

What to do before dumping a database

Before dumping a database, verify that it is working correctly. More specifically, run the following commands prior to each database dump:

- 1 dbcc checkdb.
- 2 dbcc checkalloc or dbcc checkalloc with the fix option. See “Fixing and preventing allocation errors” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for information about running these commands in multi-user mode and how to prevent false allocation errors from dbcc commands.)

Additional information

If problems persist, this procedure may not be sufficient to clean up the index corruption, and you should contact Sybase Technical Support with the information listed in “Reporting errors” on page 4, including output from dbcc checkalloc or dbcc tablealloc and dbcc checkdb.

Version in which this error is raised

All versions

Error 2626

Severity

22

Message text

Illegal attempt to insert duplicate key row in the clustered index for object '%.*s' in database '%.*s'.

Explanation

Error 2626 is raised when you attempt to insert a row into a table with a clustered index, and the key value being inserted matches an existing index key value, which was unexpected.

Error 2626 can be caused by:

- A corrupt clustered index
- An Adaptive Server problem

Action

Corrective action depends on the type of table, and whether the table is partitioned. If the problem persists after applying the correction, call Sybase Technical Support.

Tables with more than one partition

- 1 Unpartition the table:

```
1> alter table <table_name> unpartition
```

```
2> go
```

Where *<table_name>* is the object named in the message text.

2 Repartition the table:

```
1> alter table <table_name> partition
<number_of_partitions>
2> go
```

Tables with one partition

1 If the object encountering the error is a system table, perform the procedure in “Fixing a system table corrupted index” in the *Troubleshooting and Disaster Recovery* guide to repair the index. Then go to step 5.

2 If the object encountering the error is a user table, identify the clustered index on the table:

```
1> sp_helpindex <table_name>
2> go
```

Where *<table_name>* is the object named in the message text.

3 To ensure that the information needed to re-create the index is available, run `sp_helpindex` on the index prior to dropping it.

4 Drop the index, then re-create the index, or create a non-clustered index. This eliminates the 2626 error in most cases.

5 Run `dbcc checktable` on the table to verify that the corruption is gone.

Version in which this error is raised

All versions

Error 2628

Severity

21

Message text

```
In database '%.*s', page %ld is linked forward to page
%ld, but that page is linked backward to page %ld.
Please report this internal error to Sybase Technical
Support.
```

Note This error may be caused by a hardware problem.

Explanation	<p>When updating a table or index, Adaptive Server may determine that there is not enough room on a page to insert additional data. A new page is then allocated for the object. Existing rows may be moved to this page and new rows inserted. Error 2628 is raised when Adaptive Server fetches a page by following the forward pointers, examines the existing page linkage, and finds that the backward pointer does not point to the previous page as expected. Since the linkage problem could be worsened by continuing the page split, the operation is aborted.</p> <p>When dbcc detects this problem, it raises error 2503 (Table Corrupt: Page linkage is not consistent; check the following pages: (current page#=%ld; page# pointing to this page=%ld; previous page# indicated in this page=%ld)).</p>
Action	<p>Error 2628 indicates a corruption of page pointers for the object. To resolve the error:</p> <ol style="list-style-type: none">1 Identify the object using one of the pages identified in the message by performing the procedure in “Finding an object name from a page number” in the most recent <i>Troubleshooting and Disaster Recovery</i> guide. <hr/> <p>Warning! The procedure referenced in step 1 uses the dbcc page command. Use the dbcc page command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.</p> <hr/> <ol style="list-style-type: none">2 If the page is a data page (the index ID = 0 (zero)), drop and recreate the clustered index on the table.3 If the page is an index page (the index ID is greater than 0 (zero)), drop and recreate the index.4 Run dbcc pglinkage on the table to verify that the problem is corrected. Be sure to check both next page and previous page pointers.
Version in which this error is raised	All versions

Error 2630

Severity	21
Message text	In database '%1!', page %2! is linked backward to page %3!, but that page is linked forward to page %4!. Please report this internal error to Sybase Technical Support.
Explanation	In the course of updating (or deleting) data, all of the rows may be moved off a page and then this page would be deallocated. When removing a page from the page chain, Adaptive Server checks to see if the previous pointer is consistent. The 2630 error is raised when Adaptive Server discovers that the previous page pointer is pointing to a page that should have its forward pointer pointing back to the page that pointed to it. This only occurs on data-only-locked (DOL) tables.

Action To resolve the error:

- 1 Identify the object using one of the pages identified in the message by following the procedure in “Finding an object name from a page number” in the most recent *Troubleshooting and Disaster Recovery* guide.

Warning! The `dbcc` commands listed below are undocumented and nonstandard commands. The commands are provided “as is” without any warranty. Sybase Technical Support does not support these commands for general use. Although the command syntax is provided here for reference, use these commands only in the specific diagnostic situations and with the specific syntax shown in this section. Failure to do so could cause performance problems or database corruption.

For example:

```
dbcc traceon (3604)
dbcc page (<database name>, <page-number>)
```

Note In version 15.0.2 and later use:

```
set switch on 3604
```

to allow output to appear to your terminal device.

- 2 If the page is a data page (the index ID = 0(zero)), drop and recreate the clustered index on the table. Or, use `select into` a new table from this table, or `bcp out`.

If the page is an index page (the index ID is greater than 0(zero)), drop and recreate the index. Run `dbcc tablealloc` on the table to verify that the problem is corrected.

Additional information This error may be caused by a hardware problem. If 2630 errors persist, start Adaptive Server with trace flags – T12617 – T12618 – T3605 to collect additional information and if necessary, contact Sybase Technical Support.

Version in which this error is raised Version 12.5.1 and later

Error 2631

Severity 19

Message text The maximum number of index levels '%d' has been reached for index %S_PTINFO. Drop and re-create the index with a smaller key.

Explanation During an index split, error 2631 is raised if the split causes the height of the index tree to increase beyond the maximum supported height.

Action Drop and recreate the index with fewer columns included to make the index smaller. Configuring the server for a page size larger than the default 2K page size allows index keys to be much larger.

Version in which this error is raised Version 12.0.0.7 and later

Procedure Manager Errors (2800s)

This chapter contains error messages for the Adaptive Server Enterprise Procedure Manager.

Error	Page
Error 2805	373
Error 2806	374
Error 2811	375
Error 2812	376
Error 2824	378
Error 2835	379
Error 2844	380

Error 2805

Severity

20

Message text

Bad pointer 0x%lx encountered while remapping stored procedure '%.*s'. Must re-create procedure.

Explanation

When a stored procedure is created, Adaptive Server stores an efficient, normalized form of the procedure called a query tree in the `sysprocedures` table. When you upgrade the server, the query trees for stored procedures and certain other database objects may be remapped to be compatible with the new release.

Error 2805 is seen:

- When you upgrade from a prior release and there a remap failure occurs during the upgrade.
- When the tree from the `sysprocedures` table becomes corrupt.

This error may be accompanied by error 2812.

Action

Resolve the problem by dropping and re-creating the procedure.

Version in which this error is raised All versions

Error 2806

Severity 20

Message text Stored procedure '%.*s' is corrupted. Must re-create procedure.

Note This error can be caused to hardware problems.

Explanation Procedure buffers are data structures that manage Adaptive Server stored procedures (and other compiled objects like triggers, rules, defaults, check constraints, and views) in the procedure cache. The server uses one procedure buffer for every copy of a named object stored in the procedure cache.

The Procedure Buffer Manager maintains an array of procedure buffers in memory, and controls and synchronizes access to procedure buffers. It maintains a usage count to determine whether a procedure buffer is in use.

Error 2806 is raised when the Procedure Buffer Manager detects one of these problems:

- While looking for an available procedure buffer for a given stored procedure or trigger, a buffer with a usage count greater than 1 is found. Buffers for a stored procedure or a trigger are not shared and should never have a usage count that exceeds 1.
- While reading rows from `sysprocedures`, the Procedure Buffer Manager finds a header but no additional rows for the object.
- While reading rows from `sysprocedures`, the Procedure Buffer Manager finds out-of-sequence rows for the object.

Error 2806 results from corrupted procedure buffers, which can be related to hardware or operating system problems.

Action

- If this is an isolated error on a user stored procedure or trigger, drop and re-create the procedure to resolve the problem.

- If the error occurs on one or more system stored procedures (that is, procedures beginning with the `sp_` prefix), recreate the `sybsystemprocs` database where system procedures reside. See “Recovering the `sybsystemprocs` database” in the *System Administration Guide: Volume 2* for instructions.
- If other errors are reported (2806 errors during installation of a new server are often accompanied by error 605 or other errors), follow the directions in this manual for recovering from those errors.

Additional information Have the information listed in “Reporting errors” on page 4 ready if you call Sybase Technical Support.

Version in which this error is raised All versions

Error 2811

Severity 18

Message text Cannot create procedure dbid %d, objid %ld, with a group number of %d.

Explanation This error occurs when Adaptive Server is unable to insert a row into the `sysprocedures` table. This can happen during creation of a procedure, view, or trigger, or during remapping of a compiled object.

The most common reason for error 2811 is lacking enough data space to add the procedure, and it is often raised following an 1105 error.

Action Determine what errors occurred before the 2811 error by examining the:

- Server error log
- Errors the user saw when running the application
- Application log file

Additional information If error 1105 occurred, refer to “Error 1105” on page 173 for recovery instructions. If errors other than 1105 occurred, follow the recovery instructions in this manual for those errors. If this manual does not include instructions for the error that occurred, call Sybase Technical Support.

Version in which this error is raised All versions

Error 2812

Severity	16
Message text	Stored procedure '%.*s' not found. Specify owner.objectname or use sp_help to check whether the object exists (sp_help may produce lots of output).
Explanation	This error occurs when you attempt to execute a stored procedure that does not exist. If the procedure actually does exist, that is, it appears when sp_help is run with no parameters, error 2812 can occur if the procedure name was not fully qualified.

Adaptive Server's naming convention for database objects is:

```
<database>.<owner>.<object_name>.<column_name>
```

The default value for *<database>* is the current database and the default value for *<owner>* is the current user. Remember that the owner is part of the object, so two different users can have two different procedures with the same *<object_name>* in the same database, for example user1.proc1 and user2.proc1.

If the procedure is not owned by the user attempting to execute it, and the procedure is not owned by the database owner (“dbo”), then all references to it must include the owner name. For example, user1 creates a procedure called proc1. Any users other than user1 must prefix the procedure name with the owner name when executing it; that is, exec user1.proc1.

System stored procedures are an exception to the naming convention. These procedures reside in the sybssystemprocs database, are owned by the Sybase system administrator, and have names starting with “sp_”. You can execute system stored procedures from any database and they will reference the system catalog of the current database. The search path for all system stored procedures (sp_*) is the:

- 1 Current database
- 2 sybssystemprocs database
- 3 master database

If error 2812 occurs when you run the *installmaster* script, the sp_configure procedure was not found. You can ignore this.

If error 2812 occurs when you try to dump or load a database, you may be trying to dump to an Adaptive Server rather than to a Backup Server.

Action

- If you do not know who owns a procedure, use `sp_help` to display the owner. Run this procedure without any parameters to display objects owned by other users. See *Reference Manual: Procedures* for instructions.

Alternatively, use the following query to determine who owns the stored procedure you are attempting to execute:

```
1> select name,owner=user_name(uid)
2> from sysobjects
3> where name = "<procedure_name>"
4> go
```

- If the procedure does not appear in the query's output, the procedure is either in a different database or does not exist. If you do not own the procedure, qualify the procedure name with the owner name:

```
1> execute <owner_name>.<procedure_name>
2> go
```

- For procedures used by many database users, the database owner should create the procedure, which allows any user to find the procedure without specifying an owner name. If the procedure is not in the database where it is executed, fully qualify the procedure name with the database name:

```
1> exec <database_name>.<user_name>.<procedure_name>
2> go
```

The owner name is not needed if you or the database owner own the procedure:

```
1> exec <database_name>.<procedure_name>
2> go
```

Execute permission must be provided so other users can execute this procedure, but no permissions are required to see the text of the procedure.

System stored procedures

If error 2812 occurs on system stored procedures, it may be resolved by running the *installmaster* script, which installs all system procedures and initializes various other Adaptive Server structures.

Backup Server

If error 2812 occurs when you perform a dump or load, you may be trying to dump to an Adaptive Server rather than to a Backup Server. Check your Backup Server assignment:

```
1> sp_helpserver SYB_BACKUP
2> go
```

The `network_name` column is the *interfaces* file that contains the entry for this server. Check the status of the server:

```
1> exec SYB_BACKUP...sp_ps
2> go
```

Open Server applications use the stored procedure `sp_ps` while Adaptive Servers do not. Therefore, if `sp_ps` fails because it is not found, most likely the server defined for `syb_backup` is an Adaptive Server and not a Backup Server. In that case:

- 1 Find out how the Backup Server should be defined.
- 2 Drop the `syb_backup` entry using `sp_dropserver`.
- 3 Add the correct entry using `sp_addserver`.

Additional information

Refer to the *Transact-SQL User's Guide* for information about naming conventions.

Refer to the *Reference Manual: Procedures* for information about using `sp_dropserver` and `sp_addserver`.

Version in which this error is raised

All versions

Error 2824

Severity

19

Message text

```
Process %d cannot hold PROC_BUF 0x%1x named '%.*s'
because it holds another PROC_BUF 01%1x named '%.*s'. A
process can only hold one view, rule, or default at a
time.
```

Explanation

A process can hold many procedures and triggers, but it can hold only one view, rule, or default at a time. If a process already holds a view, rule, or default, and Adaptive Server tries to add another one for the process, error 2824 is raised.

Error 2824 can be caused when:

- The view referenced by stored procedures or triggers has been dropped and re-created. Stored procedures and triggers are not re-resolved when a view they reference is dropped and re-created.
- A procedure cache corruption has occurred.

Action	Use <code>sp_help view_name</code> to determine whether a view that is referenced by stored procedures or triggers has been dropped and re-created. If the view's creation time is later than the creation time of the procedure or trigger to which the view refers, drop and re-create the stored procedures and triggers that reference that view. If error 2824 persists, there is probably a procedure cache corruption. Stop and restart Adaptive Server to clear the corruption.
Version in which this error is raised	All versions

Error 2835

Severity	16
Message text	The procedure tree is old. Use the <code>sp_remap</code> procedure to remap all the procedures in this database.
Explanation	<p>When you execute a stored procedure, Adaptive Server checks to determine whether the procedure is already in cache. If the procedure is not in cache, Adaptive Server reads the procedure into the <code>sysprocedures</code> cache and recompiles it. Before Adaptive Server recompiles the procedure, it checks the Adaptive Server release number of the procedure to make sure it is up-to-date for the current release.</p> <p>Error 2835 occurs when:</p> <ul style="list-style-type: none"> • The procedure is not up-to-date for the current release. • The query remapping phase of an upgrade failed. In this case, you may need to remap all the query trees.
Action	<p>Use <code>sp_remap</code> to remap each procedure, trigger, rule, default, and view in the current database:</p> <pre>1> use <database_name> 2> go 1> sp_remap <object_name> 2> go</pre> <p>Where <code><database_name></code> is the name of the database where the object resides and <code><object_name></code> is the name of the object to remap.</p>

	If you have many objects to remap, perform the procedure in “Remapping all database objects” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Additional information	Refer to the <i>Reference Manual: Procedures</i> for information about using <code>sp_remap</code> .
Version in which this error is raised	All versions

Error 2844

Severity	20
Message text	<p>The procedure header size was expected to be '%d' memory pages instead of '%d' memory pages.</p> <p>This is an internal error. Please contact Sybase Technical Support.</p>
Explanation	A procedure header is used to keep track of stored procedures on disk. The first procedure header size should be only one page; but an attempt has been made to write a first procedure header more than one page in size. It is necessary to keep the size of the first procedure header to one page to facilitate reading it back from disk and also to allow downgrading to previous releases. The procedure has not been written to disk yet.
Action	This is an Adaptive Server internal error. Please contact Sybase Technical Support.
Additional information	Have the information listed in “Reporting errors” on page 4 and the source of the procedure that was being created available for Sybase Technical Support.
Version in which this error is raised	Version 15.0 and later

dump and load Errors (3000s - 3200s, 4200s - 4300s)

This section contains error messages raised for Adaptive Server Enterprise's dump and load commands.

Error	Page
Error 3019	382
Error 3020	383
Error 3021	384
Error 3105	385
Error 3120	385
Error 3143	386
Error 3144	387
Error 3161	387
Error 3162	388
Error 3201	388
Error 3203	389
Error 3211	391
Error 3212	391
Error 3216	392
Error 3225	393
Error 3230	393
Error 3233	394
Error 3240	395
Error 4204	396
Error 4205	400
Error 4207	401
Error 4216	402
Error 4221	403
Error 4222	404
Error 4305	406
Error 4322	406

Error 3019

Severity	10 (when sent to a client) 18 (when sent to the Adaptive Server error log)
Message text	Unable to write statistics for object %ld in database %ld. Please run update statistics on this table after loading this database or prior to accessing the table in the loaded database.

Explanation The `systabstats` system table stores table- and index-level statistics such as page and row counts, empty pages, number of leaf and non-leaf pages in an index, cluster ratios and average length of data and index rows for each object.

Descriptors are in-memory representations of Adaptive Server objects. Table descriptors contain (among other information) current statistics for the table. Adaptive Server periodically flushes these statistics to `systabstats` so that queries have access to the most current statistics for a table and its indexes. Statistics are flushed to `systabstats` during:

- checkpoints
- database dumps
- housekeeper tasks
- update statistics
- `optdiag` runs
- `sp_flushstats`

Error 3019 is raised when Adaptive Server attempts to flush table statistics to `systabstats` during a database dump, but due to contention for the `systabstats` row or some other reason, the server is unable to write out the statistics.

Error 3019 is raised with the following states:

State	Meaning
1	Cannot flush statistics during a transaction.
2	Failed attempt to flush statistics to <code>systabstats</code> .

Action This is not a serious error, but does indicate that the dump does not contain the latest statistics maintained in memory for the table. Consequently, when the dump is loaded, `systabstats` will not contain the latest statistics.

Resolve the issue by running `update statistics` for the table after loading the database.

The syntax for `update statistics` appears in the *Reference Manual: Commands*.

Version in which this error is raised Version 11.9.2 and later

Error 3020

Severity	10 (when sent to a client) 18 (when sent to the Adaptive Server error log)
Message text	Unable to write statistics for index %ld of object %ld in database %ld. Please run update statistics on this index after loading this database or prior to using this index in the loaded database.
Explanation	<p>The <code>systabstats</code> system table stores table- and index-level statistics such as page and row counts, empty pages, number of leaf and non-leaf pages in an index, cluster ratios and average length of data and index rows for each object.</p> <p>Descriptors are in-memory representations of Adaptive Server objects. Table descriptors contain (among other information) current statistics for the table. Adaptive Server periodically flushes these statistics to <code>systabstats</code> so that queries have access to the most current statistics for a table and its indexes. Statistics are flushed to <code>systabstats</code> during:</p> <ul style="list-style-type: none">• checkpoints• database dumps• housekeeper tasks• update statistics• <code>optdiag</code> runs• <code>sp_flushstats</code> <p>Error 3020 is raised when Adaptive Server is unable to flush the statistics for an index to <code>systabstats</code> during a database dump. Since the <code>systabstats</code> on the dump will not contain the latest statistics for the index, a subsequent database load will result in <code>systabstats</code> containing outdated statistics for the index.</p> <p>This is not a serious error, since the statistics can be readily updated. Queries can run normally even when <code>systabstats</code> does not contain the latest index statistics for the table.</p>
Action	Resolve the issue by running <code>update statistics</code> after loading the database:

- 1 Identify the table and object named in the message. To determine the table name, enter:

```
1> use <database_name>
2> go

1> select object_name (<object_id>)
2> go
```

Where *<object_id>* is the object identified in the message.

- 2 Run update statistics to ensure that the statistics for the index in question are included. Several options are available with update statistics, for example:

- update statistics *<table_name>* *<index_name>*
- update index statistics *<table_name>*
- update index statistics *<table_name>* *<index_name>*

The syntax for update statistics appears in the *Reference Manual: Commands*.

Version in which this error is raised

Version 11.9.2 and later

Error 3021

Severity	10
Message text	Unable to write statistics for all indexes of object %1! in database %2!. Please run update statistics on all indexes after loading this database or prior to using the indexes in the loaded database.
Explanation	This error is similar to error 3020, except that it affects all indexes for an object, not just one.
Action	Refer to the action for error 3020, following it for each index for the object.
Version in which this error is raised	Version 11.9.2 and later

Error 3105

Severity	16
Message text	Data on dump will not fit into current database. Need %d Mbyte database.
Explanation	This error occurs when an attempt is made to load a database dump into a database smaller than the dumped database. You cannot use <code>dump database</code> or <code>load database</code> to shrink a database, regardless of the amount of actual data in the database.
Action	<p>If your goal is to load the database from the dump, you must increase the size of the database being loaded into to at least the configured size of the dumped database. You can do this with either the <code>alter database</code> command or by dropping and re-creating a larger database.</p> <p>If your goal is to shrink the total size of a database, you cannot use <code>dump and load</code>. You must copy all the data out of the database using <code>bcp</code>, drop the database and re-create a smaller database, then re-create the database structure with scripts, and reload the data with <code>bcp</code>.</p>
Additional information	The database being loaded into must have identical data and log mapping, and identical segment mappings, as the dumped database. Refer to “Error 2558” on page 314 for more information.
Version in which this error is raised	All versions

Error 3120

Severity	16
Message text	The database you are attempting to LOAD was DUMPed under a different sort order ID (%d) or character set ID (%d) than the ones running on this server (SortOrd = %d, CharSet = %d). If the sort orders differ, at least one of them is non-binary.
Explanation	A database cannot be loaded into a server using <code>load database</code> unless the sort order of the server that dumped the database matches the sort order of the server into which the database is being loaded. Error 3120 occurs when you try to load a database on an Adaptive Server with a different sort order from that of the dumped database.

Action	Use one of the following options to resolve the error: <ul style="list-style-type: none">• Move the data using <code>bcp</code> rather than <code>load</code>.• Change the sort order of the Adaptive Server. Refer to “Configuring Character Sets, Sort Orders, and Languages” in the <i>System Administration Guide: Volume 1</i> for instructions.• If the server into which you want to load the database is a new server:<ol style="list-style-type: none">a Reinstall the server with the correct sort order.b Re-create the databases with the <code>load</code> option.c Retry the load.
Version in which this error is raised	All versions

Error 3143

Severity	16
Message text	The master database you are attempting to load was dumped under a different sort order id (%d) or character set id (%d) than the one running on this server (sort order id = %d, character set id = %d).
Explanation	Either the character set or the sort order (or both) of the master database dump differs from the current values in the Adaptive Server. It is possible that the current character set is not even available in <code>syscharsets</code> in the master database being loaded, so this must be disallowed.
Action	Change the character set or sort order of the running server to match the master database being loaded. If necessary, use trace flag 3100 to allow the load to continue. <hr/> Warning! Use this trace flag only after you verify that the server character set or sort order exists in <code>syscharsets</code> of the master database being loaded. <hr/>
Version in which this error is raised	Version 11.5.1.1 and later

Error 3144

Severity	16
Message text	<pre>Since this load may render the master database unusable, this load will only be permitted under the 3100 traceflag. Please contact Sybase Technical Support regarding the correct usage of this traceflag.</pre>
Explanation	<p>This message is raised when there is a mismatch of either character set or sort order between the server configuration and the values in the master database dump. It is possible that the current server character set is not even available in <code>syscharsets</code> in the master database being loaded, so this must be disallowed. This message always appear after error 3143 is raised.</p>
Action	<p>If it is not possible to change the character set and sort order of the running server to match the master database dump, use the trace flag 3100 to allow the dump to be loaded.</p> <hr/> <p>Warning! Use this flag only after you verify that the server character set or sort order exists in <code>syscharsets</code> of the master database dump being loaded. If you use trace flag 3100 to force the load of a master database dump without proper entries in <code>syscharsets</code>, the master database will be unusable, and must be recreated.</p> <hr/>
Version in which this error is raised	Version 11.5.1.1 and later

Error 3161

Severity	17
Message text	<pre>LOAD DATABASE for database `%.*s` failed during cross- platform conversion: insufficient memory to allocate context structure.</pre>
Explanation	<p>Error 3161 is raised when there is not enough memory to allocate an internal structure needed to hold information for cross platform database translation and cleanup.</p>

Action	Raise the value of the configuration parameter procedure cache size and retry the load.
Version in which this error is raised	Version 12.5.3 and later

Error 3162

Severity	17
Message text	<code>LOAD DATABASE for database `%.*s` failed during cross-platform conversion: insufficient memory to allocate translation dictionary.</code>
Explanation	Error 3162 is raised when there is not enough memory to allocate an internal structure needed for the translation dictionary for cross platform database translation.
Action	Raise the value of the configuration parameter procedure cache size and retry the load.
Version in which this error is raised	Version 12.5.3 and later

Error 3201

Severity	16
Message text	<code>No dump device has been specified.</code>

Note This error can be caused by a hardware problem.

Explanation This error occurs when Adaptive Server is unable to access a dump device during a database dump. This error is fatal and stops the dump from completing.

Error 3201 is usually caused by one of the following conditions:

- The device you specified for the dump is offline or otherwise unavailable to Adaptive Server.

- Permissions for the dump device are not set correctly for user “sybase” or the user performing the dump (read and write privileges are required).
- A previous dump aborted and Adaptive Server believes the dump device is still in use.

Action

- 1 Verify that the device you specified is defined for your Adaptive Server:

```
1> select * from master..sysdevices
2> where status = 16 or status = 24
3> go
```

low	high	status	cntrltype	name	phyname	mirrorname
0	20000	16	2	tapedump1	/dev/rmt4	NULL
0	20000	16	3	tapedump2	/dev/rst0	NULL

If the device is not there, use `sp_addumpdevice` to add the device.

- 2 At the operating system level, check the permissions for the dump device and for user “sybase” or the user performing the dump (read and write privileges are required).
- 3 If you still have not resolved the error, check the operating system error log for operating system errors. If no errors are logged in the operating system error log and there are no other errors in the Adaptive Server or Backup Server error log, contact Sybase Technical Support for assistance.

Additional information

Refer to “Specifying the Database and Dump Device” in the *System Administration Guide: Volume 2* for information.

If you call Technical Support, have the information on hand that is listed in “Reporting errors” on page 4, including the output from `select * from sysdevices`.

Version in which this error is raised

All versions

Error 3203

Severity

16

Message text	Backup Server has detected a SQL Server error.
	<hr/> Note This error can be caused by a hardware problem. <hr/>
Explanation	<p>Error 3203 is raised when a problem occurs with a remote procedure call to the Backup Server. In this case, the error is on the Adaptive Server side.</p> <p>The error message output includes:</p> <ul style="list-style-type: none">• <i>vsn</i> – the virtual socket number.• <i>return</i> – the return value: 0 means successful; -2 means failure.• <i>status</i> – the ending status, displayed in decimal value. The most common value is 524288 which usually means an I/O error. <p>Possible causes for this error are:</p> <ul style="list-style-type: none">• Write protect is turned “on” (it must be “off” when loading a tape because the device is opened using read and write permissions).• No dump exists on the media being accessed; for example, you are using a blank tape or specified the wrong device.• The media that contains the dump is not readable.
Action	<ul style="list-style-type: none">• Check the Adaptive Server and Backup Server error logs to determine the cause of the error being sent from the Adaptive Server.• Test the connection between the Backup Server and the Adaptive Server. Log into the Adaptive Server using <code>isql</code> and type:<pre>1> execute <backupserver>...sp_ps 2> go</pre>Where <i><backupserver></i> is the name of your Backup Server. This command executes <code>sp_ps</code> on the Backup Server.
Additional information	If you need to call Sybase Technical Support, have the information on hand that is listed in “Reporting errors” on page 4.
Version in which this error is raised	All versions

Error 3211

Severity	20						
Message text	DUMP failed in database %.*s: error occurred executing RPC.						
Explanation	<p>When a dump database command is processed, the Backup Server scans for allocated pages to determine what pages to dump. It begins this scan as soon as it receives an RPC (remote procedure call). Error 3211 is raised when an Adaptive Server RPC to the Backup Server returns an error.</p> <p>Error 3211 occurs with the following states:</p> <table border="1"> <thead> <tr> <th>State</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>An error occurred in the RPC stream.</td> </tr> <tr> <td>2</td> <td>The RPC itself returned an error.</td> </tr> </tbody> </table>	State	Meaning	1	An error occurred in the RPC stream.	2	The RPC itself returned an error.
State	Meaning						
1	An error occurred in the RPC stream.						
2	The RPC itself returned an error.						
Action	<ul style="list-style-type: none"> Check the Adaptive Server and Backup Server error logs to determine the cause of the error being sent. Test the connection between the Backup Server and the Adaptive Server. Log into the Adaptive Server using <code>isql</code> and type: <pre>1> execute <backupserver>...sp_ps 2> go</pre> <p>Where <code><backupserver></code> is the name of your Backup Server. This command <code>sp_ps</code> on the Backup Server.</p> 						
Additional information	If you need to call Technical Support, have the information on hand that is listed in “Reporting errors” on page 4.						
Version in which this error is raised	All versions						

Error 3212

Severity	20
Message text	DUMP failed in database %.*s: unexpected status executing RPC: %d.

Explanation	Dumps and loads are performed by an Open Server program known as Backup Server. Adaptive Server communicates with Backup Server using remote procedure calls (RPCs). When a <code>dump database</code> command is processed, for example, Adaptive Server uses RPCs to send parameters to Backup Server. Error 3212 is raised when an Adaptive Server RPC to the Backup Server returns an unexpected status that Adaptive Server cannot process.
Action	<ul style="list-style-type: none">• Check the Adaptive Server and Backup Server error logs to determine the cause of the error.• Test the connection between the Backup Server and the Adaptive Server. Log into the Adaptive Server using <code>isql</code> and type:<pre>1> execute <backupserver>...sp_ps 2> go</pre>Where <code><backupserver></code> is the name of your Backup Server. This executes <code>sp_ps</code> on the Backup Server.
Additional information	If you need to call Sybase Technical Support, have the information on hand that is listed in “Reporting errors” on page 4.
Version in which this error is raised	All versions

Error 3216

Severity	10
Message text	Backup Server session id is: %ld. Use this value when executing the 'sp_volchanged' system stored procedure after fulfilling any volume change request from the Backup Server.
Explanation	Dumps and loads are performed by an Open Server program known as Backup Server. During the backup and restore process, it may be necessary to change tape volumes. If Backup Server detects the need to change the currently mounted volume, it requests a volume change by sending messages to either the client or its operator console. After mounting another volume, the operator notifies Backup Server by executing the <code>sp_volchanged</code> system procedure on Adaptive Server.

Error 3216 is an informational message that Backup Server issues when starting a dump or load operation. The message notifies the operator of the current session ID, which the operator needs if it becomes necessary to execute `sp_volchanged` during the dump or load operation.

Action This is an informational message. No action is necessary.

Version in which this error is raised All versions

Error 3225

Severity 21

Message text I/O failed while attempting to clear pages starting at page %ld of disk number %ld.

Note This error can be caused by a hardware problem.

Explanation During a load database, create database, or alter database command, Adaptive Server initializes the non-allocated database pages. Error 3225 occurs if an I/O error occurs during this process.

This error can be caused by an operating system problem such as insufficient disk space for a file system or a disk going offline.

Action Check your operating system error logs to determine whether you have run out of disk space or are having other problems that could explain the I/O errors. Refer to “Checking the operating system error log” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for assistance.

Version in which this error is raised All versions

Error 3230

Severity 24

Message text For logical disk device '%.*s', neither primary device

```
'%.*s' nor secondary device '%.*s' is active. %s cannot
continue.
```

Explanation

During a dump or load, Adaptive Server must translate the logical device names into physical device names and pass these to Backup Server, which performs the dump or load. For each device, Adaptive Server first checks whether the primary physical device is in the active state; if not, it checks whether the secondary (mirror) device is active.

Error 3230 is raised when neither the primary nor secondary devices are active.

Warning! This is a fatal error that breaks your connection to Adaptive Server.

The error may be caused by:

- The device being offline
- An Adaptive Server problem, which results when you create a device using a previously used `vdevno`.

Action

Examine the Adaptive Server error log and the operating system log to see if the devices used in the dump or load are offline or otherwise unavailable to the server. See “Checking the operating system error log” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for assistance.

If the devices are active, shut down and restart Adaptive Server, and retry the dump or load operation.

Version in which this error is raised

All versions

Error 3233

Severity

17

Message text

```
%s for database '%S_DBID' failed: an alarm function
could not be installed.
```

Explanation

Dumps and loads are performed by an Open Server program named Backup Server. In preparation for the dump, Adaptive Server must install an alarm function, which is a timing mechanism required during the dump. Using the alarm function, Backup Server goes into a “sleep” state at an appropriate point and awakens periodically to check the progress of the dump.

Error 3233 is raised when you perform a dump, but there are not enough alarms available to allow Adaptive Server to install an alarm function for the dump, causing the dump to abort.

Action

Available alarms are controlled by the configuration parameter `number of alarms`. To resolve the problem, change the value of `number of alarms` to a number larger than the current value.

Refer to “No more alarms available” in the Configuration Parameters chapter of the *System Administration Guide: Volume 1* for directions on increasing the number of alarms.

Because the `waitfor` command uses alarms, frequent use of this command can also lead to 3233 errors. Check your use of this command in statement blocks, stored procedures, and transactions.

Version in which this error is raised

All versions

Error 3240

Severity

18

Message text

Unable to obtain the SQL Server's current working directory. The SQL Server must have search permission on its current working directory and the name must be shorter than %d bytes.

Explanation

Dumps and loads are performed by an Open Server program known as Backup Server. When you issue a `dump database` or `load database` command, Adaptive Server communicates certain information about its working environment to Backup Server. Error 3240 is raised when Adaptive Server is unable to determine its current working directory for Backup Server.

Action

This error commonly occurs when Adaptive Server is started from a directory other than the `sybase home` directory; the user ID of the user starting Adaptive Server may not have the appropriate permissions for the directory from which it was started.

To resolve the error, check that the `SYBASE` variable is set to the correct `sybase home` directory and is no longer than 255 bytes.

Make corrections if necessary, and restart Adaptive Server.

Additional information	Scripts that start up Adaptive Server must always set the directory to the sybase home directory before starting Adaptive Server.
Version in which this error is raised	All versions

Error 4204

Severity	17
Message text	Unable to continue logged version of DUMP TRAN. No space left in database; use NO_LOG option or run ALTER DATABASE. Else, an old transaction may be active; see Syslogshold table.
Explanation	This error occurs when Adaptive Server runs out of space in the transaction log during a dump transaction command.
Action	Use any of the following options to recover from error 4204: <ul style="list-style-type: none">• If any old transactions are active, determine whether they can be terminated. If such transactions do exist and they can be terminated, dump the transaction log, which may free up database space.• Perform a dump transaction with the <code>no_log</code> option. This command removes the inactive part of the log without making a backup copy and without recording the procedure in the transaction log (performs a checkpoint).• Extend the transaction log segment.

Each option is described below.

Checking for old transactions

In the `master` database, there is a table called `syslogshold`. This table contains information about each database's oldest active transaction (if any) and Replication Server truncation point (if any) for the transaction log. This table is built dynamically when you query it.

Check `syslogshold` for old transactions for the database for which the error occurred:

```
1> use master
2> go
```



```
1> select * from syslogshold
2> where dbid = <database_ID>
3> go
```

Determine whether the oldest active transaction can be terminated (perhaps it was left active unintentionally). If you can terminate this transaction and then execute a dump transaction with `no_log` command, space may be returned to the database. Continue this process until there are no other old transactions that can be terminated. If, after terminating some old transactions, you still get error 4204, try one of the other options described in this write-up.

Recovering using *dump transaction with no_log*

- 1 Use the following method to determine approximately how many pages the transaction log occupies. Enter the commands that correspond to the version of Adaptive Server Enterprise you are using:

- Versions earlier than 15.0:

```
1> use <database_name>
2> go
```

```
1> select data_pgs(8,doampg)
2> from sysindexes where id=8
3> go
```

- Version 15.0 and later:

```
1> select data_pages (<database id>,8)
2> go
```

Where “8” is the ID for `syslogs`. The result is the number of data pages (2K pages on most operating systems) that the transaction log occupies.

Note The query results may be inaccurate by as many as 16 pages, but using the query is much faster than counting the rows in the transaction log.

- 2 Dump the transaction log using the `no_log` option.
- 3 Repeat step 1. If the transaction log now occupies significantly fewer pages, continue with the next step now.

If, however, there are still a large number of pages in the `syslogs` table, an outstanding transaction is probably preventing the log from being cleared. If this is the case, do one of the following:

- Follow the instructions in “Detecting and clearing long-running transactions” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.
 - Restart Adaptive Server and repeat step 2. When Adaptive Server starts and this database is recovered, the outstanding transaction is rolled back, allowing the log to be cleared by a subsequent dump transaction.
- 4 Dump the database now, because the `no_log` option prevents subsequent changes recorded in the log from being used to recover from a media failure. You must perform a `dump database` now to ensure the recoverability of subsequent database modifications using transaction logs.

Note This database dump is not required if your backup and recovery procedures do not include saving the transaction logs for media failure recovery.

If you are concerned that your transaction log is not large enough for your Adaptive Server, refer to information about estimating the transaction log size within “Creating and Managing User Databases” in the *System Administration Guide: Volume 2* before deciding to increase its size. If you choose to increase the size, refer to *Reference Manual: Commands* for more information about the `alter database` command used to increase log size. Remember that once they have been increased, log and database sizes cannot be decreased.

Extending the transaction log segment

Warning! If you use this option to extend the transaction log segment, you will not be able to get that space back later.

- 1 Refer to “Estimating the Transaction Log Size” in the chapter “Creating and Managing User Databases” in the *System Administration Guide: Volume 2* before deciding whether to increase the transaction log size.
- 2 To get information about the log segment, enter commands like the following example and replace “`test`” with the name of the database in the error message:

```
1> use test
2> go
```

```

1> sp_helpsegment logsegment
2> go

```

segment	name	status
2	logsegment	0

device	size	free_pages
dev1	1.0MB	512

table_name	index_name	indid
syslogs	syslogs	0

- 3 If you have more space in the current log device fragment, increase the amount of space allocated for the transaction log with commands like the following:

```

1> use master
2> go

1> alter database <database_name>
2> log on <device_name> = <additional_space>
3> go

```

Where:

- *<database_name>* is the name of the database for which the log device fragment is to be extended.
- *<device_name>* is the name of the database device on which to locate the database extension.
- *<additional_space>* is the number of megabytes to extend the log device fragment.

- 4 If you do not have any more space in the current log segment, you can use `sp_extendsegment` to extend the range of the log segment to another database device (the database must already exist on that device fragment before you run these commands):

```

1> use <database_name>
2> go

1> sp_extendsegment logsegment, <database_name>,
2> <device_name>
3> go

```

Where:

- `<database_name>` is the name of the database where the segment is to be extended.
- `<device_name>` is the name of the additional database device for the log segment.

Warning! If you use `sp_extendsegment` to extend the range of the log segment to another database device, and the other device contains data, the segment on which that data exists will become a log segment. This situation can cause you to run out of space in the database earlier than expected and may cause other problems, including 2558 errors.

Version in which this error is raised

All versions

Error 4205

Severity

16

Message text

Syslogs does not exist in its own segment in database '%S_DBID' with segmap '%ld' with logical start page number of '%ld'. You cannot use DUMP TRANSACTION in this case, use DUMP DATABASE instead.

Explanation

The dump transaction command copies the transaction log, which resides in the syslogs table, recording any database changes made since the last database dump or transaction log dump. Error 4205 is raised when you attempt dump transaction, but the log and data share the same device. Transaction log dumps are only possible when data and log reside on separate devices.

Since transaction logs are used for up-to-date recovery in the event of a disk failure, they are useful only if kept separate from the data device.

Action

Use dump database. This will make a copy of the entire database, including both data and the transaction log.

Additional information

Unless you are creating small, noncritical databases, always place the log on a separate device. This allows you to dump transaction logs, ensures full recovery from hard disk crashes and has many other benefits including performance improvement. For details, refer to “Placing the Transaction Log on a Separate Device” in the chapter “Creating and Managing User Databases” in the *System Administration Guide: Volume 2*.

When data and log share device space, it is possible to move the transaction log to its own device by using the system stored procedure `sp_logdevice`. For related information, see the write-up for “Error 4222” on page 404 in this manual.

Version in which this error is raised

All versions

Error 4207

Severity

16 State 2

Message text

Dump transaction is not allowed because a non-logged operation was performed on the database. Dump your database or use dump transaction with `truncate_only` until you can dump your database.

Explanation

This error occurs when you have performed a minimally logged operation such as `bcp in` or `select into` on the database, followed by a dump transaction to a dump device. The dump transaction is not allowed because your database contains data that cannot be recovered from the log. The log could not be used to recover the database properly. For instance, if a user does a `select into` followed by an update, data for the `select into` is not recorded in the log, so the update cannot be recovered from the log.

Note A minimally logged operation is one in which changes to the data itself are not logged to the transaction log. All changes to the allocation structures resulting from the operation are logged to the transaction log.

Note that simply enabling the `select into / bulkcopy` option does not, in itself, cause this condition. The option must be set and a minimally logged operation executed before this error message is reported during a dump transaction to a dump device.

Dumping the database creates a new save point on which subsequent transaction logs are based. Therefore, dumping the database ensures that the minimally logged changes are recoverable and dumping the transaction log to a device is once again allowed.

Action	<p>Perform a dump database. Disabling the <code>select into/bulkcopy</code> option using <code>sp_dboption</code> will not allow you to dump the log. You must dump the database to clear the 4207 error condition. You may want to disable the <code>select into/bulkcopy</code> option to prevent any further nonlogged operations in your database.</p> <p>You may need to use <code>dump transaction with truncate_only</code> until you can perform a full database dump. If you must do this, you will lose the ability to recover up-to-the-minute changes in data in case of a media failure.</p>
Version in which this error is raised	All versions

Error 4216

Severity	21
Message text	<code>%s failed in database %.*s: unexpected end of log encountered on page %ld while scanning for page %ld.</code>

Note This error may be caused by a hardware problem.

Explanation

The transaction log consists of a series of log pages chained together in a doubly linked list; each log page has a page header with pointers to the next and previous log pages. The previous page pointer for the first page is zero, as is the next page pointer for the last page, indicating the beginning and end of the log chain respectively.

Error 4216 is raised when Adaptive Server finds what it considers to be a premature end of the transaction log chain during execution of a command. This error is raised during `dump transaction` or `dump database` and is often seen in conjunction with error 4219.

Error 4216 can be caused by a bad page chain or allocation (`syslogs` corruption).

Action	<ol style="list-style-type: none">Determine the status of <code>syslogs</code>:<pre>1> use <database_name> 2> go 1> dbcc checktable (syslogs) 2> go</pre>
--------	--

```
1> dbcc tablealloc (syslogs)
2> go
```

Where *<database_name>* is the name of the database appearing in the error message.

dbcc reports any problems found in syslogs, and the number of data pages being used by the log. If your log is on a separate device, dbcc checktable also tells you how much space is used and how much is free.

- 2 If no problems were reported, the transaction log is clean. Call Sybase Technical Support for assistance.
- 3 If problems were reported, try to fix them using the sections in this manual that describe those errors. If errors are still reported, call Sybase Technical Support.

Additional information

For information about estimating and managing the transaction log size, refer to “Creating and Managing User Databases” in the *System Administration Guide: Volume 2*.

Before calling Technical Support, have the information listed in “Reporting errors” on page 4 on hand, including the output of dbcc checktable and dbcc tablealloc.

Version in which this error is raised

Version 11.0 and later

Error 4221

Severity

16

Message text

DUMP TRANSACTION to a dump device is not allowed where a truncate-only transaction dump has been performed after the last DUMP DATABASE. Use DUMP DATABASE instead.

Explanation

The dump database command makes a copy of the entire database, including both the data and the transaction log. dump database does not truncate the log.

The dump transaction command copies the transaction log, recording any database changes made since the last database or transaction log dump. After copying the log, it truncates the inactive portion. The dump transaction with truncate_only command, on the other hand, removes transactions from the log without making a backup copy.

Error 4221 occurs in the following situations:

- If you attempt a dump transaction command, and the most recent dump consisted of a dump transaction with `truncate_only`. This is not allowed because the dump transaction would produce a transaction dump that could not be applied, since the prior dump transaction with `truncate_only` removed transactions upon which the current transactions may depend. To ensure recoverability, you must dump the database with `dump database` each time it is necessary for you to run a dump transaction with `truncate_only` command.
- If you use the `log on clause to create database` to place the transaction log on a separate database device, and you attempt a dump transaction prior to dumping the database.

Action

Perform a `dump database` first to back up the database and log. Follow this up with the appropriate dump transaction command.

Make sure the `trunc log on chkpt database` option is turned “off.” When this option is turned “on,” you cannot dump the transaction log.

Version in which this error is raised

All versions

Error 4222

Severity

16

Message text

DUMP TRANSACTION for database '%S_DBID' failed: log page (%ld) is on a data-only segment. This is probably due to an incomplete move of the log to a device separate from the data. Complete the log migration or contact the System Administrator.

Explanation

The dump transaction command copies the transaction log, recording any database changes made since the last database dump or transaction log dump. After copying, it truncates the inactive portion of the log. Transaction log dumps are only possible when data and log reside on separate devices.

If data and log share device space, it is possible to move the transaction log to its own device by using the system stored procedure `sp_logdevice`.

Error 4222 is raised when you execute `dump transaction`, but the first page of the transaction log is found to be on a data-only segment of the database. Since the log is not isolated from the data, the transaction dump fails.

The likely cause of this error is that the steps to move the log to a separate device were initiated, but were not complete at the time you issued `dump transaction`. To understand this, note that `sp_logdevice` moves future allocation for the transaction log to the new device. However the log remains on the original device until the extent that is currently in use has been filled and the transaction log has been dumped. Error 4222 indicates that the log in its current state cannot be migrated to its own device.

Action

To resolve the error:

- 1 Run `dbcc checktable (syslogs)` to determine which page is currently in use.
- 2 Execute enough transactions to fill the extent currently in use. The following commands will fill at least one extent on the log (8 pages):

```
1> create table dummy1 (c1 char(255), c2 char(255))
2> go
```

```
1> insert dummy1 values ("a", "a")
2> go 16
```

```
1> drop table dummy1
2> go
```

- 3 Confirm that there are no currently active transactions on the database device, then run `dump transaction with truncate_only`.

Note If an active transaction causes a `begin tran` to be written to the log pages in question, a transaction dump will not be able to truncate those pages. This is why you should avoid having activity on the device while you migrate the log off the data segment.

- 4 Run the system procedure `sp_helplog` to ensure that the complete log is on the new device.

Additional information

Refer to “Moving the transaction log to another device” in the chapter “Creating and Managing User Databases” in *System Administration Guide: Volume 2* for details on moving the transaction log.

Version in which this error is raised

All versions

Error 4305

Severity	16
Message text	Specified file '%.*s' is out of sequence. Current time stamp is %S_DATE while dump was from %S_DATE.
Explanation	<p>During a load transaction, Adaptive Server checks the timestamp on the dumped database and each dumped transaction log to determine whether the sequence is correct. If the load sequence of the transaction logs is incorrect, error 4305 occurs and the load fails.</p> <p>If you are missing a transaction log dump, you can only restore up to the transaction log dump before the missing dump.</p> <p>In addition, you cannot load multiple transaction log dumps that have been created with the <code>no_truncate</code> option. The <code>no_truncate</code> option assumes that nothing further will be done with the transaction log after that point: it is only useful for up-to-the-minute recovery of a database after failure of the data device. If you have used the <code>no_truncate</code> option and your database is still okay, do a dump database as soon as possible.</p> <p>This error is serious as the load processing stops and the load transaction fails.</p>
Action	Load the log backups in the same sequence in which they were made. Check your transaction log dumps to determine whether one was loaded out of order. If you find the correct dump, load it and continue with the rest of the dumps, in order.
Additional information	For more information, refer to “load transaction” in the <i>Reference Manual: Commands</i> .
Version in which this error is raised	All versions

Error 4322

Severity	16
Message text	SQL Server cannot load a dump of the master database that was dumped from an older version of SQL Server. The current master database upgrade version is %ld, the dump master database upgrade version is %ld.

Explanation Error 4322 is raised when you try to load a dump of the master database that was created on an earlier version of Adaptive Server than your currently installed version.

Action If you know you do not have a dump of master from your current version level because you did not dump master after you upgraded, and your master database is corrupted, call Sybase Technical Support.

Otherwise, check the current upgrade version of the master dump you are trying to load. For example:

```
1> load database master from 'master.dump' with headeronly
2> go
```

```
Backup Server session id is: 54. Use this value when executing the
'sp_volchanged' system stored procedure after fulfilling any volume change
request from the Backup Server. Backup Server: 6.28.1.1: Dumpfile name
'master072670B065 ' section number 1 mounted on disk file
'/work/dputz/ase1502/ASE-15_0/install/master.dump' This is a database dump of
database ID 1, name 'master', from Sep 24 2007 12:32PM. ASE version: Adaptive
Server Enterprise/15.0.2/EBF 14328/P/Sun_svr4/OS 5.8/ase1502/2486/64-
bit/FBO/Thu May 24 12:. Backup Server version: Backup Server/15.0/EBF 13443
ESD#2/Sun_svr4/OS 5.8/ase15x/2603/32-bit/OPT/Fri Mar 24 03:05:01 2006. Database
page size is 2048. Database contains 6656 pages; checkpoint RID=(Rid pageid =
0x100c; row num = 0x13); next object ID=924527296; sort order ID=50, status=0;
charset ID=1. Database log version=7; database upgrade version=35. segmap:
0x00000007 lstart=0 vstart=[vpgdevno=0 vpvpn=4] lsize=6656 unrsvd=2982
```

The log version and upgrade version have these values:

Version	Database log version	Database upgrade version
12.0 through 12.0.0.8	5	25
12.5 through 12.5.4	6	31
15.0 and later	7	35

Additional information If you call Sybase Technical Support, have the information on hand that is listed in “Reporting errors” on page 4, including the output from the load database...with headeronly command if you have what you thought was a current dump, and the output from sp_configure "upgrade version".

Version in which this error is raised All versions

This section contains error messages for Adaptive Server Enterprise commit and abort transaction operations.

Error	Page
Error 3301	409
Error 3307	410
Error 3321	411
Error 3322	412

Error 3301

Severity

21

Message text

Invalid log record found in syslogs (logop %d).

Note This error can be caused by a hardware problem.

Explanation

The `syslogs` table is the Adaptive Server transaction log that contains a record of changes made in the database and in the order in which they occur. Adaptive Server scans the records in the transaction log during database recovery and other activities. Error 3301 occurs if Adaptive Server encounters an invalid record while scanning the transaction log. `logop` refers to the type of record (log operation) encountered. For example, a `logop` of 18 is logged by a `save` transaction command.

Error 3301 is raised with the following states:

State	Meaning
1	An unexpected log record was encountered during recovery.
9	An unexpected log record was encountered while translating log records associated with an obsolete log format (for a previous Adaptive Server version) to the format for the current Server version.
10	An out-of-range <code>logop</code> number was encountered while translating log records associated with an obsolete log format (for a previous Adaptive Server version) to the format for the current Server version.

State	Meaning
11	An invalid log record was found in the the analysis, undo, or redo log operations handling.
12	An invalid log record was found while running the analysis pass of recovery.
13	An update log record marked part of a twin did not have either an undo-only or redo-only status set.
14	An out-of-range <code>LOGOP</code> number was encountered while regenerating a distribution map from the <code>XDISTMAP</code> transaction log records.

	<p>This is a serious error that indicates a problem with the transaction log. Corrective action depends on the source of the error.</p>
Action	<p>A 3301 error during recovery (boot time, load database, or load transaction recovery) marks the database suspect. The database is identified in the error log. Restore the database from a good backup as follows:</p> <ol style="list-style-type: none"> 1 Drop the database. If the drop fails, follow the instructions in “Dropping a database when drop database fails” in Chapter 4 of the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide. 2 Create a database for loading. Verify that the database you create is at least as large as the original database, and that its <code>sysusages</code> values match the original values. For more information, refer to “create database” in the <i>Reference Manual: Commands</i>. 3 Load the database from backup. (Refer to “load database” in the <i>Reference Manual: Commands</i>.) 4 Use the online database command to make the database available for use. <p>A 3301 error during rollback of a transaction will generally result in a “phantom” process that continues to hold locks that in time may block other transactions that require the same locks. Contact Sybase Technical Support for assistance.</p>
Additional information	<p>Refer to “Developing a Backup and Recovery Plan” in the <i>System Administration Guide: Volume 2</i> for instructions about how to safely create, dump, load, and re-create databases.</p>
Version in which this error is raised	<p>All versions</p>

Error 3307

Severity

21

Message text	Process %d was expected to hold logical lock on page %ld in database %S_DBID for object %S_OBJID, index number %d.
Explanation	This error occurs when rolling back a transaction that made changes to non-leaf pages of an all-pages-locked (APL) index and the expected locks for this operation were not already held.
Action	<p>Look for other errors in the Adaptive Server error log as well as in your operating system error log to find out the specific source of the problem and clear those errors first, as they might be the actual cause of the error.</p> <p>This error can result in locks being held by the process that rolled back. Restarting Adaptive Server releases any locks that this transaction might still hold. However, this does not clear the root cause of this error.</p> <p>If there are no other errors in the Adaptive Server error log or in your operating system error log, contact Sybase Technical Support for assistance.</p>
Additional information	<p>For more information about transactions, refer to “transaction” in the <i>Transact-SQL User's Guide</i>.</p> <p>Before calling Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the output of sp_lock and sp_who before restarting Adaptive Server.</p>
Version in which this error is raised	All versions

Error 3321

Severity	21
Message text	Redo of CREATE INDEX in transaction ID (%d,%d) for table name '%.*s', object ID '%d', failed to find sort partition log record. Recovery of database '%.*s' cannot continue.
Explanation	<p>This error is raised during a roll forward of a create index in load transaction. While scanning the log, expected log records were not found.</p> <p>Error 3321 is raised with these states:</p>

State	Meaning
1	Unable to find the log record that describes the properties of the index that is being created.

State	Meaning
2	Unable to find all the log records for every partition in the table.
3	Unable to find the log record that describes the input table for which an index is being created.

Action	Load the database and transaction dumps that occurred up to the transaction dump that caused the problem, then recreate the indicated index. This transaction dump, and any subsequent transaction dumps, cannot be loaded. Please contact Sybase Technical Support for any necessary help in recovering the database.
Additional information	See the most recent version of <i>Troubleshooting and Disaster Recovery</i> for instructions on database recovery.
Version in which this error is raised	Version 15.0 and later

Error 3322

Severity	21
Message text	In redo of sort, the number of partitions expected (%d) does not equal the number of partitions logged (%d) by transaction ID (%d,%d) operating on table '%.*s', object ID '%.*s'. Recovery of database '%.*s' cannot continue.
Explanation	For a roll forward of a create index in load transaction, the actual number of partitions for which an index is to be created does not match the expected number of partitions.
Action	Load the database and transaction dumps that occurred up to the transaction dump causing the problem, then recreate the indicated index. This transaction dump, and any subsequent transaction dumps, cannot be loaded. Please contact Sybase Technical Support for any necessary help in recovering the database.
Additional information	See the most recent version of <i>Troubleshooting and Disaster Recovery</i> for information about database recovery.
Version in which this error is raised	Version 15.0 and later

Recovery Errors (3400s, 12500s)

This section contains error messages for Adaptive Server Enterprise database recovery. Database recovery occurs during Adaptive Server start-up, load database, and load transaction.

Error	Page
Error 3401	414
Error 3403	414
Error 3404	416
Error 3414	416
Error 3429	417
Error 3434	419
Error 3445	420
Error 3446	421
Error 3447	422
Error 3452	423
Error 3454	424
Error 3462	425
Error 3470	426
Error 12546	427
Error 12547	428
Error 12551	429
Error 12583	429
Error 12587	430

Error 3401

Severity

21

Message text

Rec_logbounds: getnext SCAN_RID of last checkpoint failed on Rid from sysdatabases. %S_RID.

Explanation When you restart Adaptive Server, recovery scans the log to find the last checkpoint record to determine the active portion of the log. Error 3401 is raised when recovery is unable to find a checkpoint record.

Error 3401 occurs with the following states:

State	Meaning
1	A log scan was started, but an attempt to find the checkpoint record in the log failed.
3	An attempt to scan the log to find the checkpoint log record failed.

Action

Good backup available

If you have a good backup, restore the database as follows:

- 1 Drop the database. If the drop fails, follow the instructions in “Dropping a database when drop database fails” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide.
- 2 Create a database for load. Make sure the database you create is at least as large as the original database (and that its `sysusages` values match the original values). For more information about creating a database for load, refer to “create database” in the *Reference Manual: Commands*.
- 3 Load the database from backup. (Refer to “load database” in the *Reference Manual: Commands*.)
- 4 Use the online database command to make the database available for use.

No good backup available

If you do not have a good backup, contact Sybase Technical Support. They may be able to help you recover from this error.

Additional information

Have the information on hand that is listed in “Reporting errors” on page 4 when you call Sybase Technical Support:

Version in which this error is raised

All versions

Error 3403

Severity

22

Message text

During recovery initialization, page %ld was

encountered. This page belongs to object %ld, not the log.

Note This error may be caused by a hardware problem.

Explanation During recovery, Adaptive Server scans the transaction log until the last page is found. During this scan, if a page is found that has an object ID on the page header that is not the same as that of `syslogs`, error 3403 occurs.

This may occur either as a result of a hardware failure or due to an Adaptive Server problem.

Warning! This is a serious error and you must rebuild the affected database either using `bcp` or from clean backups.

Action If you have a good backup, restore the database from your backup:

- 1 Drop the database. If the drop fails, follow the instructions in “Dropping a database when drop database fails” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide.
- 2 Create a database for load. Make sure the database you create has sizes at least as large as those in `sysusages` for the original database (and that all other `sysusages` values match the original values). For more information about creating a database for load, refer to “create database” in the *Reference Manual: Commands*.
- 3 Load the database from backup. (Refer to “load database” in the *Reference Manual: Commands*.)
- 4 Use the online database command to make the database available for use.

To prevent this error from being raised in the future, either perform a normal shutdown of Adaptive Server (specifically, do not use the `nowait` option), or checkpoint each database that is being used before shutting down Adaptive Server.

Additional information Refer to “Developing a Backup and Recovery Plan” in the *System Administration Guide: Volume 2* for complete information about how to safely create, dump, load, and re-create databases.

Version in which this error is raised All versions

Error 3404

Severity	21
Message text	<code>Rec_complete: Could not open controlling database (id %d) of controlling database in multi-db transaction.</code>
Explanation	<p>A multi-database transaction spans more than one database in a server. The coordinating database is the one in which the transaction started. The participant databases are other databases involved in the multi-database transaction. A two-phase protocol is used to commit this transaction. When Adaptive Server is restarted after an unexpected shutdown, the recovery of a participant database in a multi-database transaction may need to open the coordinating database to determine the outcome of the transaction if it is in a prepare-to-commit state. The 3404 error is raised if Adaptive Server is unable to open the coordinating database.</p> <p>Specifically, error 3404 results when Adaptive Server is attempting to recover a transaction that started in a user database but modified data in <code>master</code> database. When this happens, Adaptive Server is unable to access the transaction log for the user database since the device on which the log resides is not yet open. Consequently <code>master</code> cannot be recovered. This error is frequently accompanied by errors 913 (see “Error 913” on page 152) and 3414 (the next error write-up).</p>
Action	<p>Use one of these options to correct this problem:</p> <ul style="list-style-type: none">• Restore the database specified in error 3404 from a clean backup.• Contact Sybase Technical Support. They may be able to help you re-create the database in question. <p>To avoid a recurrence of this error, do not start transactions in user databases that modify tables in <code>master</code>.</p>
Additional information	Before calling Technical Support, have the information available that is listed in “Reporting errors” on page 4.
Version in which this error is raised	All versions

Error 3414

Severity	10
----------	----

Message text	<p>Versions earlier than 15.0:</p> <p style="padding-left: 40px;">Database '%.*s' (dbid %d): Recovery failed. Check the SQL Server errorlog for further information as to the cause.</p> <p>Version 15.0 and later:</p> <p style="padding-left: 40px;">Database '%.*s' (dbid %d): Recovery failed. Check the ASE errorlog for further information as to the cause.</p>
Explanation	<p>This error occurs during startup when Adaptive Server could not complete the recovery of the database listed in the error message.</p> <p>In most cases, you cannot use the database until whatever caused the error has been corrected because Adaptive Server marks the database suspect. Some commands (like drop database) are allowed to execute.</p> <p>Error 926 is a related error which is raised when attempting to access a database that could not recover. Refer to the write-up for “Error 926” on page 160.</p>
Action	<p>To determine why recovery failed, examine the Adaptive Server error log for any errors prior to the 3414 error. It is important to identify the errors before the first occurrence of the 3414 error because subsequent attempts to start Adaptive Server may not give the detailed error information you need to diagnose the problem.</p> <p>If you do not have sufficient information to recover from the previous errors, you can recover from known, clean backups or contact Sybase Technical Support for assistance.</p>
Version in which this error is raised	All versions

Error 3429

Severity	21
Message text	Error recovering database '%.*s' - could not connect to commit service to check completion status of xact:

	%S_RID.
Explanation	<p>Adaptive Server enables cross-server transactions using either an external transaction coordinator (like Microsoft's Distributed Transaction Coordinator) or through a CT-Library two-phase commit application programming interface (API) called the Sybase two-phase commit service. In both cases, a two-phase commit protocol is used to commit the cross-server transaction.</p> <p>When you restart Adaptive Server after an unexpected shutdown, the recovery of this transaction might find the transaction in a prepared-to-commit state. The coordinator must be contacted to determine the final transaction status of the transaction. The 3429 error is raised when the external transaction coordinator (DTC) or commit service (Sybase two-phase commit) cannot be contacted.</p> <p>If the transaction is coordinated by DTC, consult the Microsoft documentation to determine the cause of the error.</p> <p>If the transaction is using the Sybase two-phase commit service, an internal probe utility is used to query the commit server. As a standard client, probe requires:</p> <ul style="list-style-type: none"> • A valid interface must exist for the commit server. • The commit service must be running and not in single-user mode. • The probe version must be the same as the Adaptive Server version. <p>In most cases, recovery will mark the database suspect after raising the 3429 error.</p>
Action	<ol style="list-style-type: none"> 1 Determine what is preventing Adaptive Server from contacting either the coordinator or the commit server. 2 To remove the suspect state from the database, refer to "Resetting a database's "suspect" status" in the chapter "Other Useful Tasks" in the most recent version of <i>Troubleshooting and Disaster Recovery</i>. <hr/> <p>Note This should only be done if it is acceptable that the database be made available without having been fully recovered.</p> <hr/> <ol style="list-style-type: none"> 3 Shut down and restart Adaptive Server to recover the database.
Additional information	Refer to the <i>Open Client DB-Library/C Reference Manual</i> for information about the two-phase commit service.
Version in which this error is raised	All versions

Error 3434

Severity	20
Message text	Cannot change sortorder. Server shutting down. Restart to continue with sortorder unchanged.
Explanation	<p>When you restart a server after changing the server's sort order, the recovery process must rebuild system tables in both <code>master</code> database and user databases, and rebuild indexes that are affected by the sort order change. Error 3434 occurs when recovery is unable to recreate a system index or table following the sort order change.</p> <p>The error can be raised:</p> <ol style="list-style-type: none"> 1 When you change from a case-sensitive sort order to a case-insensitive sort order, and conflicts arise between key values that are differentiated only by case. For example, key values 'Joe' and 'joe' will cause a conflict when rebuilding the index with a case-insensitive sort order. <ul style="list-style-type: none"> Error 3436 ("Cannot rebuild index %d for the '%.*s' table in the '%.*s' database.") precedes error 3434 in this situation. 2 When one or more databases for load exist in the server. You cannot change sort order when the server contains a database that was created for load. <p>Error 3434 rolls back the transaction in progress, reverts to the original sort order, and shuts down the server.</p>
Action	<p>Check the error log to obtain more information about the error. Actions depend on the likely cause of the error:</p> <ol style="list-style-type: none"> 1 If the error is related to key value conflicts, determine the table and index from the 3434 error message, and check for "duplicate key" messages in the error log. Correct the reported conflict and retry the sort order change. 2 If the server contains a database created for loading a dump (for load database), take the following steps to correct the problem: <ol style="list-style-type: none"> a Drop the for load databases. b Manually change the sort order on the server. See "Manually changing sort order or default character set" in the chapter "Server Recovery Tasks in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide. c Restart the server.

After changing the sort order, refer to “Reconfiguring the character set, sort order, or message language” in the chapter “Configuring Character Sets, Sort Orders, and Languages” in the *System Administration Guide: Volume 1* and perform the steps described there to guarantee the integrity of your data.

Additional information

- You can achieve similar results without changing the sort order to case-insensitive by using the upper or lower functions in your queries. For example, if `tableX` has `col1` with values “sybase”, “Sybase”, and “SYBASE”, you could run the following `isql` statement:

```
select cols from tableX where lower(col1) = "sybase"
```

This treats all values as lowercase for the query.

- In many cases, you cannot reload your data from a database dump after reconfiguring the sort order. Refer to “Reconfiguring the character set, sort order, or message language” in the *System Administration Guide: Volume 1* for details.

Version in which this error is raised

All versions

Error 3445

Severity

10

Message text

Versions earlier than 15.0:

```
SQL Server could not bring database '%.*s' online.
```

Version 15.0 and later:

```
ASE could not bring database '%.*s' online.
```

Explanation

This error occurs when Adaptive Server cannot bring a database online for one of the following reasons:

- An attempt to update the database log version fails. Adaptive Server tries to update the log version if it is earlier than the current Adaptive Server log version.
- Adaptive Server failed to clear the offline status bit after an upgrade.

Action

Look at the error messages that precede the 3445 message to determine why the 3445 error occurred, and resolve those problems.

If you cannot solve the problem:

- 1 Determine the current log version for the database. For example, to determine the log version of a database called `test_db`, use the following commands:

```
1> load database testdb from '/tmp/test.dump' with headeronly
2> go
```

Backup Server session id is: 6. Use this value when executing the ``sp_volchanged'` system stored procedure after fulfilling any volume change request from the Backup Server.

Backup Server: 6.28.1.1: Dumpfile name `testdb0726708DD9 ' section number 0001 mounted on disk file `~/tmp/test.dump'

This is a database dump of database ID 4, name 'testdb', from Sep 24 2007 10:05AM. ASE version: Adaptive Server Enterprise/15.0.2/EBF 14328/P/Sun_svr4/OS 5.8/ase1502/2486/64-bit/FBO/Thu May 24 12:.. Backup Server version: Backup Server/15.0/EBF 13443 ESD#2/Sun_svr4/OS 5.8/ase15x/2603/32-bit/OPT/Fri Mar 24 03:05:01 2006. Database page size is 2048. Database contains 15360 pages; checkpoint RID=(Rid pageid = 0x50f; row num = 0x18); next object ID=800002850; sort order ID=50, status=0; charset ID=1.'Database log version=7; database upgrade version=35.' 'segmap: 0x00000007 lstart=0 vstart=[vpgdevno=0 vpvpn=11780] lsize=15360unrsvd=14272

- 2 Call Sybase Technical Support. Have the information ready that is listed in “Reporting errors” on page 4, including the output from step 1.

Version in which this error is raised

All versions

Error 3446

Severity 16

Message text You do not have privilege to bring database '%.*s' online.

Explanation The online database command marks a database available for public use after a normal load sequence and, if needed, upgrades a loaded user database and transaction log dumps to the current version of Adaptive Server.

Only a System Administrator, Database Owner, or Operator can execute online database. Error 3446 occurs when you do not have any of these privileges and you attempt to execute an online database command.

Action	Make sure you have the needed privileges or ask a user who does have the needed privileges to run the online database command.
Additional information	Refer to the <i>Reference Manual: Commands</i> for information about using the grant and online database. Refer to “Granting and revoking roles” in the <i>System Administration Guide: Volume 1</i> for information about using the “sa” account and roles in Adaptive Server.
Version in which this error is raised	All versions

Error 3447

Severity	10
Message text	<p>Versions earlier than 15.0:</p> <pre>Database `%.*s' appears to be in the process of being loaded; SQL Server will not bring it online automatically. Use the ONLINE DATABASE command to bring this database online.</pre> <p>Version 15.0 and later:</p> <pre>Database '%.*s' appears to be in the process of being loaded; ASE will not bring it online automatically. Use the ONLINE DATABASE command to bring this database online.</pre>
Explanation	<p>During a load database, Adaptive Server takes the database being loaded offline. When the load is complete, you bring the database back online manually with the online database command. If your backup strategy involves loading a sequence of transaction logs after a load database, the fact that the database will be unavailable for use until you enter the online database command will allow you to complete your loads without interference from other processes changing that database.</p> <p>Error 3447 occurs during the recovery phase of Adaptive Server start-up, when the recovery process attempts to bring all the recovered databases online automatically and Adaptive Server sees that the database named in the error message is being loaded.</p> <p>This error does not occur for master, as Adaptive Server always brings master online.</p>

Action	Finish the load database sequence and then use online database to make the database available for use.
Additional information	Refer to the <i>Reference Manual: Commands</i> for information about using the online database command.
Version in which this error is raised	All versions

Error 3452

Severity	20
Message text	Database '%.*s': upgrade item %d depends on item %d, which could not be installed. Please refer to previous messages for the cause of the failure, correct the problem and try again.
Explanation	<p>Databases can be upgraded during an <code>online database</code> command or when you are using the <code>sybinit</code> upgrade utility. Many upgrade steps depend on other upgrade steps having been done previously.</p> <p>Error 3452 occurs when the first upgrade step shown in the message depended on the second reported step, and that second step has failed. The database upgrade does not complete when this error occurs.</p> <p>You will not be able to bring the database named in the error message online until you resolve the problem that led to the 3452 error.</p>
Action	<ol style="list-style-type: none"> 1 Find the previous message describing the failure of the dependent (second) step. 2 Fix the problem that caused the failure. If appropriate, follow the directions for a failed upgrade in the Adaptive Server installation and configuration guide. 3 If you were running <code>online database</code> when error 3452 occurred, issue the <code>online database</code> command to restart the database upgrade.
Version in which this error is raised	All versions

Error 3454

Severity 10

Message text Versions earlier than 15.0:

```
Database '%.*s': SQL Server could not completely upgrade
this database; upgrade item %d could not be installed.
```

Version 15.0 and later:

```
Database '%.*s': ASE could not completely upgrade this
database; upgrade item %d could not be installed.
```

Explanation

Databases can be upgraded during an online database command or when you are using the `sybinit` upgrade utility. Error 3454 is an informational message telling you that Adaptive Server could not completely upgrade the database listed in the message. The 3454 error is preceded by other errors that show why the upgrade failed.

You cannot bring the database named in the error message online until you resolve the problem that led to the 3454 error.

In the following example, these messages appeared during start-up, when Adaptive Server attempted to bring a database online but failed to upgrade it:

```
00:2006/11/21 18:34:36.51 server Database 'test' appears to be at an older
revision than the present installation; SQL Server will assess it, and upgrade
it as required.
```

```
00:2006/11/21 18:34:36.59 server Error: 1105, Severity: 17, State: 1
00:2006/11/21 18:34:36.61 server Can't allocate space for object
'sysattributes' in database 'test' because the 'system' segment is full. If you
ran out of space in syslogs, dump the transaction log. Otherwise, use ALTER
DATABASE or sp_extendsegment to increase the size of the segment.
```

```
00:2006/11/21 18:34:36.63 server Error: 3460, Severity: 20, State: 1
00:2006/11/21 18:34:36.65 server Database 'test': upgrade could not record the
installation of upgrade item '80'. Please refer to previous error messages to
determine the problem. Fix the problem, then try again.
```

```
00:2006/11/21 18:34:36.68 server Error: 3451, Severity: 20, State: 1
00:2006/11/21 18:34:36.70 server Database 'test': upgrade has failed for this
database. Please refer to previous messages for the cause of the failure,
correct the problem and try again.
```

```
00:2006/11/21 18:34:36.73 server Error: 3454, Severity: 20, State: 1
00:2006/11/21 18:34:36.75 server Database 'test': SQL Server could not
completely upgrade this database; upgrade item 80 could not be installed.
```

```
00:2006/11/21 18:34:36.76 server SQL Server could not bring database 'test'
```

online.

Action	<ol style="list-style-type: none"> 1 Look at the error messages preceding the 3454 error message to determine what happened. 2 Refer to the recovery steps suggested for those errors. 3 Correct the problem. 4 If the error occurred when you were running online database, issue the online database command for that database after the problem is fixed. Adaptive Server will automatically try to finish the upgrade. If the error occurred during an Adaptive Server upgrade, follow the recovery directions in the Adaptive Server installation and configuration guide for your platform.
Version in which this error is raised	All versions

Error 3462

Severity	10
Message text	Database '%.*s': upgrade could not install optional upgrade item '%d'. Please refer to previous error messages to determine the problem. Upgrade will continue.
Explanation	<p>An optional upgrade item is part of Adaptive Server's upgrade process, but is not critical for a successful upgrade. Error 3462 means that the upgrade task did not complete, but the problem did not prevent the upgrade process from continuing.</p> <p>This error is preceded by other messages related to the upgrade problem.</p>
Action	Check for other upgrade messages just prior to the 3462 error. After upgrade completes, correct the problem using the write-ups for those messages that are listed in this guide.

If you are upgrading, and the upgrade process encounters duplicate rows in the `syskeys` system table, the 3462 error may be accompanied by error 2615 (see “Error 2615” on page 361). This happens because the upgrade exposes a previously hidden column, `spare1`, setting its contents to zero. If the database previously contained two or more `syskeys` rows that were identical except for this column, attempting to set `syskeys.spare1` to zero will produce 2615 errors, accompanied by 3462 errors.

Note The duplicate `syskeys` entries are not caused by the upgrade process; rather, the upgrade process reveals such entries if they exist.

If you see 2615 errors during upgrade:

- 1 Make sure that you are upgrading using one of the recommended migration paths.
- 2 After upgrade completes, you still must correct the problem by removing duplicate `syskeys` rows. Follow the steps for resolving error 2615 in “Case 2” on page 363.

Additional information

Call Sybase Technical Support. Have the information ready that is listed in “Reporting errors” on page 4.

Version in which this error is raised

All versions

Error 3470

Severity

10

Message text

Versions earlier than 15.0:

```
SQL Server could not completely upgrade database '%.*s',
but the database was online when upgrade began, so it
will be left online.
```

Version 15.0 and later:

```
ASE could not completely upgrade database '%.*s', but
the database was online when upgrade began, so it will
be left online.
```

Explanation	When Adaptive Server tries to bring an already online database to the online state and it finds that some elements of the database have not been upgraded, it tries to complete the database upgrade. Error 3470 occurs when that attempt to upgrade the database fails.
Action	<p>1 Determine the current upgrade version for the database. For example, to determine the upgrade version of a database called <code>test_db</code>, use the following commands:</p> <pre>1> load database testdb from 'test.dump' with headeronly 2> go</pre> <p>Backup Server session id is: 48. Use this value when executing the 'sp_volchanged' system stored procedure after fulfilling any volume change request from the Backup Server. Backup Server: 6.28.1.1: Dumpfile name 'testdb0726708DD9 ' section number 1 mounted on disk file '/tmp/testdb.out' This is a database dump of database ID 4, name 'testdb', from Sep 24 2007 10:05AM. ASE version: Adaptive Server Enterprise/15.0.2/EBF 14328/P/Sun_svr4/OS 5.8/ase1502/2486/64-bit/FBO/Thu May 24 12:. Backup Server version: Backup Server/15.0/EBF 13443 ESD#2/Sun_svr4/OS 5.8/ase15x/2603/32-bit/OPT/Fri Mar 24 03:05:01 2006. Database page size is 2048. Database contains 15360 pages; checkpoint RID=(Rid pageid = 0x50f; row num = 0x18); next object ID=800002850; sort order ID=50, status=0; charset ID=1. ''Database log version=7; database upgrade version=35.''segmap: 0x00000007 lstart=0 vstart=[vpgdevno=0 vpvpn=11780] lsize=15360 unrsvd=14272 <p>2 Call Sybase Technical Support. Have the information ready that is listed in “Reporting errors” on page 4, including the output from step 1 and <code>sp_configure "upgrade version"</code>.</p> </p>
Version in which this error is raised	All versions

Error 12546

Severity	18
Message text	<p>Versions earlier than 15.0:</p> <p>During undo the timestamp to be assigned to a page is less than the current page timestamp. Page = ld, page timestamp = (%04x %08lx), new page timestamp = (%04x %08lx).</p>

Version 15.0 and later:

```
During undo the timestamp to be assigned to a page is
less than the current page timestamp. Page = %ld, object
id = %ld, page timestamp = (%04x %08lx), new page
timestamp = (%04x %08lx).
```

Explanation	During the process of an undo operation, Adaptive Server attempts to verify that the timestamp to be assigned to a page is greater than the timestamp that is currently on the page. If not the timestamp to be assigned is not greater, error 12546 is raised. This is particularly important during reboot and online undo because the database timestamp is derived from the log.
Action	This is a nonrecoverable error. Either restore the database from backup, or call Sybase Technical Support for suggestions on what can be done to bypass recovery on this database.
Additional information	Contact Sybase Technical Support for instructions on running the dbcc log command to identify the incorrect log record.
Version in which this error is raised	Version 12.0.0.7 and later

Error 12547

Severity	18
Message text	Warning: Memory Allocation to store recovery information failed.
Explanation	During the process of database recovery or putting a database online, Adaptive Server stores information about the databases in a temporary structure in procedure cache. Adaptive Server was unable to allocate sufficient memory for this operation.
Action	Configure additional procedure cache.
Additional information	In some cases this error just avoids that the databases cannot be recovered in parallel or in the desired order. In such cases it is not a fatal error, however in some other cases, the error will mean that the databases cannot be recovered and it is a fatal error. If this error occurs while recovering the master database, the server will be shutdown. Instead, if the error occurs while recovering a user database, the database will be marked as suspect.
Version in which this error is raised	Version 12.5.1 and later

Error 12551

Severity	16
Message text	Warning: Failed to start a recovery task.
Explanation	Adaptive Server failed to initialize a parallel recovery task. The recovery will continue with the available parallel recovery tasks. If there are no parallel recovery tasks, Adaptive Server will recover the databases in serial mode.
Action	The current start up of the Adaptive Server recovery will proceed, however Adaptive Server may have problems spawning other tasks. It may be necessary for you to increase the configuration parameter <code>number of user connections</code> to make sure sufficient resources are available to the server.
Additional information	If the 12551 error persists, it may be necessary to gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.5.1 and later

Error 12583

Severity	21
Message text	<p>Versions earlier than 15.0:</p> <pre>Unable to log a record due to internal error %1!. This process will retry at intervals of one minute.</pre> <p>Version 15.0 and later:</p> <pre>Unable to log a record in database '%.*s' (id %d) due to internal error %d. This process will retry at intervals of one minute.</pre>
Explanation	A fatal error was encountered while writing to the transaction log. This is typically seen during a failed undo operation. Rather than shutting down, the Adaptive Server will retry the operation a minute later in anticipation that what prevented the log operation from completing the first time might now be resolved.
Action	If the log operation cannot be completed after several tries, a reboot of Adaptive Server may be necessary.

Additional information	This error rarely occurs. However, if this error is seen after you restart Adaptive Server, a configurable shared memory dump (CSMD) configured for error 12583 (or module 12500) may assist Technical Support in diagnosing the cause. Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.5.2 and later

Error 12587

Severity	20
Message text	<pre>Initialization of sysroles and sysusers in database %1! failed due to missing system role '%2!' in master..syssrvroles.</pre>
Explanation	While in the process of upgrading the server or upgrading a database, Adaptive Server detected that a system role is missing.
Action	Check the <code>master..syssrvroles</code> table to see if the reported missing role is there. If this role had been manually removed by deleting a row from the <code>syssrvroles</code> table, it must be readded to the <code>syssrvroles</code> table.
Version in which this error is raised	Version 15.0 and later

Exception Handling Errors (3600s)

This section contains error messages for Adaptive Server Enterprise error and exception handling.

Error	Page
Error 3604	431
Error 3621	432
Error 3624	432
Error 3626	433
Error 3641	435
Error 3642	436
Error 3643	436
Error 3644	437
Error 3646	437
Error 3648	438

Error 3604

Severity	10
Message text	Duplicate key was ignored.
Explanation	<p>This error occurs when you try to insert duplicate key values in a column or columns which have a unique index that was created with the option <code>ignore_dup_key</code>.</p> <p>Adaptive Server ignores the commands in a transaction which try to insert a duplicate index. The rest of the commands in the transaction will complete normally.</p>
Action	<p>No action is necessary.</p> <p>If you want the row inserted:</p> <ul style="list-style-type: none"> Eliminate the unique index

or:

- Remove duplicate values. If the duplicate values are a result of incorrect data, change the data and rerun the transaction.

Additional information

Refer to “create index” in the *Reference Manual: Commands*.

Version in which this error is raised

All versions

Error 3621

Severity

10

Message text

Command has been aborted.

Explanation

This error occurs when Adaptive Server is unable to complete a command. Error 3621 may be the last message an application displays even though another message was sent by Adaptive Server when the error actually occurred. The rest of the commands in the transaction will complete normally.

Some situations that may cause error 3621 include:

- Division by 0 (zero) in an application
- Entry of a duplicate row during use of an application
- Other illegal functions in an application

Action

Look for an error before the 3621 error occurs. Resolve the problem that is causing the error before the 3621 error.

If client applications receive this error, ask your Adaptive Server system administrator to check the Adaptive Server error log.

Version in which this error is raised

All versions

Error 3624

Severity

10

Message text

Truncation error occurred.

Explanation	<p>The precision of a datatype denotes the maximum number of digits allowed in columns of that datatype. For the exact numeric types <code>dec</code>, <code>decimal</code> and <code>numeric</code>, the allowed range of data depends on the column's precision as well as its scale; that is, the maximum number of digits that are allowed to the right of the decimal point.</p> <p>Error 3624 occurs when an operation inserts data into a target column or variable but the precision or scale of the target are too small for the data. The operation fails and the command aborts.</p>
Action	Examine the data involved in the arithmetic operation to ensure that the target column or variable is sized appropriately.
Additional information	<p>The set <code>arithabort</code> and set <code>arithignore</code> options specify behavior following loss of precision or (for exact numeric types) loss of scale. For more information see:</p> <ul style="list-style-type: none"> • “Arithmetic Errors” in the <i>Transact-SQL User's Guide</i> • The Adaptive Server Enterprise reference manuals
Version in which this error is raised	All versions

Error 3626

Severity	17
Message text	<pre>The transaction was aborted because it used too much stack space. Either use sp_configure to increase the stack size, or break the query into smaller pieces. spid: %d, suid: %d, hostname: %.*s, application name: %.*s</pre>
Explanation	<p>At start-up, Adaptive Server allocates one stack for each configured user connection. These stacks are in contiguous areas of memory, with a guard area at the end of each stack. Adaptive Server periodically checks to determine whether the stack point for a user connection has entered the stack guard area adjacent to that connection's stack. Error 3626 occurs when this has happened.</p> <p>Because error 3626 occurs before any stack corruption, Adaptive Server does not shut down. Instead, it aborts the user process and rolls back any transactions that existed for that process.</p> <p>An Adaptive Server process usually encounters this error for one of two reasons:</p>

- The stack size is too small. The information placed on the stack is valid but its total size exceeds the specified stack size. This is often caused by a single, complex query which requires more than the average amount of stack space. Long *where* clauses, long *select* lists, and deeply nested stored procedures can contribute to this type of stack overflow and error 3626.
- An incorrectly handled error results in recursive calls which eventually fill the stack, regardless of its size.

If this error occurs frequently, there is increased risk that a process may overflow its stack area and corrupt another process' stack. This situation will cause Adaptive Server to shut down with this error:

```
kernel: *** Stack guardword corrupted.
```

See "Stack guardword corrupted" on page 686 for more information.

Action

If this error occurs with a transaction that contains a *where* or *where-in* clause, you may need to rewrite the query.

If this error occurs because a complex query has run out of stack area, you can resolve it in either of two ways:

- Break up the query into smaller pieces;

or

- Use this query to increase the stack size:

```
1> sp_configure "stack size", <new_value>
2> go
```

Shut down and restart Adaptive Server to implement these changes.

When changing the stack size of Adaptive Server, consider:

- The stack size should always be an even increment of 2K (2,048 bytes). If a specified stack size is not divisible by 2K, *sp_configure* rounds the value up to the next 2K increment.
- Increase the stack size in small increments to avoid making it too large. If the stack size is too large, you will lose cache space and you may be unable to start Adaptive Server.
- Adaptive Server requires more memory for each configured user connection and device when the stack size increases. Therefore, increasing the stack size without increasing the *total memory* (pre-12.5) or *max memory* (version 12.5 and later) configuration parameter will result in less total cache space for Adaptive Server, which can affect performance.

Additional information	<p>Refer to “sp_configure” in the <i>Reference Manual: Procedures</i> and “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for information about sp_configure and stack size and stack guard size configuration parameters.</p> <p>Refer to the Adaptive Server Enterprise performance and tuning guides and “Configuring Memory” in the <i>System Administration Guide: Volume 2</i> for details about configuring Adaptive Server memory.</p>
Version in which this error is raised	All versions

Error 3641

Severity	17
Message text	Failed to set alarm. Please try increasing the number of alarms in the configuration file.
Explanation	<p>When this error is raised with states 1 or 2, then the asynchronous logging service (ALS) threads are being started. Adaptive Server sets up an alarm for the ALS service to be woken up on periodic basis to check for work that need to be done. The error is raised if all alarm structures are in use at the same time and free alarm structures are not available to set the alarm. For more information on ALS, please refer to the <i>Performance and Tuning: Basics</i> section “Asynchronous log service” in the chapter “Indexing for Performance.”</p> <p>If this error is raised with state 3, then there were not enough alarm structures to set a timer for the license heartbeat.</p>
Action	<p>Regardless of the state in which the error was raised, review and increase the configuration parameter number of alarms.</p> <p>Alarms are used by a variety of processes:</p> <ul style="list-style-type: none"> • A user process executed a waitfor delay command • RepAgent threads • Checkpoints • Housekeeper tasks

Additional information	<p>The availability of alarm structures can change, depending on the server's activity. If the error persists, despite increasing the configured number of alarms, or you do not believe the Adaptive Server has too many concurrent processes that requiring alarms, capture the following outputs close to the occurrence of error 3641 and contact Sybase Technical Support:</p> <ul style="list-style-type: none">• <code>sp_who</code>• <code>sp_configure</code> number of alarms
Version in which this error is raised	Version 12.5.1 or later

Error 3642

Severity	17
Message text	<code>Unable to start flusher thread.</code>
Explanation	Adaptive Server attempted to start a process to handle log flushing for the asynchronous logging service (ALS). This attempt failed. For more information about ALS, see the <i>Performance and Tuning: Locking</i> section "Asynchronous log service."
Action	Adaptive Server uses a process slot for the thread. Verify that the number of user connections is configured sufficiently high to allow for this. Increasing number of user connections can help resolve this error.
Additional information	If the 3642 error persists after configuration changes, it may be necessary to gather a configurable shared memory dump (CSMD). See the section "Configuring a shared memory dump on conditions" in the most recent version of the <i>Troubleshooting And Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.5.1 and later

Error 3643

Severity	17
Message text	<code>Unable to start log writer thread.</code>

Explanation	Adaptive Server attempted to start a process to handle writing the log for the asynchronous logging service (ALS). This attempt failed. For more info on ALS, refer to the <i>Performance and Tuning: Locking</i> section “Asynchronous log service.”
Action	Adaptive Server uses a process slot for the thread. Check to verify that the number of user connections is configured sufficiently high to allow for this. Increasing the number of user connections can help to resolve this issue.
Additional information	If the 3643 error persists after configuration changes, it may be necessary to gather a configurable shared memory dump (CSMD). See the section “Configuring a shared memory dump on conditions” in the most recent version of the <i>Troubleshooting And Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.5.1 and later

Error 3644

Severity	17
Message text	Unable to start Asynchronous logging service.
Explanation	Adaptive Server attempted to start the asynchronous logging service (ALS) and encountered an error. There will be a prior message in the Adaptive Server error log indicating what problem was encountered.
Action	Check the error log for errors 3641, 3642, or 3643. Refer to the write-ups for those errors for the actions to perform to resolve the errors.
Additional information	For more information about ALS, see the <i>Performance and Tuning: Locking</i> section titled “Asynchronous log service.”
Version in which this error is raised	Version 12.5.1 and later

Error 3646

Severity	17
Message text	FATAL Error: Asynchronous logging server is stopped. You may have to terminate the user connections which may be

hanging.

Explanation	If the asynchronous logging service (ALS) process gets an infected stack trace, the process stops. The user connections that are waiting to be served by the ALS process may be hanging. Data modification activities are not allowed in the database. Error 3646 is raised.
Action	Kill the user connections. If the connections are not killed, restart Adaptive Server, which will roll back the incomplete database transactions. Please contact Sybase Technical Support to report this problem.
Additional information	If this error persists, then a configurable shared memory dump (CSMD) should be configured, which can be helpful in determining the root cause of this error. See the section “Configuring a shared memory dump on conditions” in the most recent version of the <i>Troubleshooting And Disaster Recovery</i> guide for instructions.
Version in which this error is raised	Version 12.5.1 and later

Error 3648

Severity	17
Message text	<pre>Async log service threads encountered a FATAL error while serving database '%.*s'. Database Log is locked and will not be released. Please look in the errorlog for more information or contact Sybase Technical Support.</pre>
Explanation	If the asynchronous log service (ALS) thread gets infected while doing a flush of the user log cache, the transaction log for that database may be inconsistent. All data change activities are disallowed in the database. The database noted in error 3648 is possibly in an inconsistent state. Processing in other databases using ALS are also stopped, because the spinlock on the transaction log is not released.
Action	Restart Adaptive Server to roll back incomplete transactions and return the database to a consistent state.
Additional information	If this error persists, then a configurable shared memory dump (CSMD) should be configured, which can be helpful in determining the root cause of this error. See the section “Configuring a shared memory dump on conditions” in the <i>Troubleshooting And Disaster Recovery</i> guide for instructions.

Version in which this
error is raised

Version 12.5.1 and later

This section contains error messages for Adaptive Server drop commands.

Error	Page
Error 3701	441
Error 3702	442
Error 3703	443

Error 3701

Severity	11
Message text	Cannot drop the %S_MSG '%.*s', because it doesn't exist in the system catalogs.
Explanation	<p>This error occurs when you try to drop an object that is not found in at least one system table where Adaptive Server expected to find it.</p> <p>Error 3701 can occur due to the following circumstances:</p> <ul style="list-style-type: none"> • The object you are trying to drop does not exist • Inconsistent system catalog tables • An Adaptive Server problem occurred
Action	<p>Make sure you entered the object name correctly in your drop command.</p> <p>If you are entering the object name correctly and the drop fails with error 3701, try to re-create the object. If your <code>create</code> command displays the following message, then your system tables are incorrect with regard to this object.</p> <pre>Msg 2714, Level 16, State 1: Line 1: There is already an object named 'object_name' in the database.</pre>

	<p>If this occurs, run <code>dbcc checkcatalog</code> and <code>dbcc checkdb</code>. Then call Sybase Technical Support to help you delete the object that is causing the error. However, because other objects may reference that object, deleting the object cleanly may be difficult. If this is the case, you may need to recover from backups.</p>
Additional information	<p>Before calling Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the output of <code>dbcc checkdb</code> and <code>dbcc checkcatalog</code>.</p>
Version in which this error is raised	<p>All versions</p>

Error 3702

Severity	16
Message text	<pre>Cannot drop the %S_MSG '%.*s', because it is currently in use.</pre>
Explanation	<p>This error occurs when Adaptive Server tries to drop a database, user table, procedure, rule, default, trigger, or view that is currently in use.</p> <p>Error 3702 can occur during:</p> <ul style="list-style-type: none"> • A drop command such as <code>drop database</code>, <code>drop table</code>, <code>drop procedure</code>, or <code>drop rule</code> • A <code>dbcc dbrepair (database_name, dropdb)</code> command <p>Error 3702 can occur for the following reasons:</p> <ul style="list-style-type: none"> • The object you are trying to drop is in use • The process using the object was abnormally terminated • Data corruption has occurred • An Adaptive Server problem occurred
Action	<p>Use <code>sp_who</code> to determine whether the database or object name in your drop command is currently in use. If it is in use, contact the user shown in the <code>sp_who</code> output.</p>

If the process using the object was abnormally terminated (for example, with an operating system kill command), the object may not have been completely removed from the Adaptive Server system tables. The process would remain invisible while keeping the object open and in use. This prevents the drop command from succeeding. In this case, shutting down and restarting Adaptive Server should clear the 3702 error. After you restart Adaptive Server, run `dbcc checkcatalog` to determine whether corruption has occurred.

If you are trying to drop a database that contains keys currently used to encrypt columns in other databases, the command will fail. Before dropping the database containing the encryption keys you must first remove the encryption or drop the database containing the encrypted columns.

If error 3702 persists, call Sybase Technical Support.

Additional information Before calling Sybase Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the output of `dbcc checkdb`, `dbcc checkcatalog`, and `dbcc checkalloc` for the database that contains the object named in the error message.

Version in which this error is raised All versions

Error 3703

Severity 11

Message text Cannot drop the %S_MSG with object-id %ld in database %d, because it doesn't exist in the system catalogs.

Explanation This error occurs when Adaptive Server fails to drop a database, user table, procedure, rule, default, trigger, or view because the object being dropped does not exist in the appropriate system table.

Error 3703 can occur during a drop command such as drop database, drop table, drop procedure, or drop rule.

Error 3703 occurs with these states.

State	Meaning
2	During a drop index command, if Adaptive Server fails to find the object ID of the table for which the index is being dropped in <code>sysobjects</code> , error 3703 occurs with state 2.
4	When Adaptive Server fails to find the referenced table in <code>sysobjects</code> for a trigger that is being dropped, error 3703 occurs with state 4.

State	Meaning
5	During any drop command, if Adaptive Server fails to find a <code>sysobjects</code> entry for any triggers belonging to the object being dropped, error 3703 occurs with state 5.
6	During any drop command, if Adaptive Server fails to find a <code>sysobjects</code> entry for the object being dropped, error 3703 occurs with state 6.
7	When Adaptive Server fails to find the referenced table in <code>sysobjects</code> for a trigger that is being dropped, error 3703 occurs with state 7.
8	When Adaptive Server fails to drop a database, user table, procedure, rule, default, trigger, or view that is not some other object's subsidiary part (for example, dropping a trigger because its associated table is being dropped), and the object ID is not found in <code>sysobjects</code> , error 3703 occurs with state 8.
9	If the object being dropped cannot be located in <code>sysobjects</code> , error 3703 occurs with state 9.
10	If the object being dropped has referential dependencies, Adaptive Server tries to delete the necessary rows from <code>sysconstraints</code> . If the ID of the reference is not found in <code>sysconstraints</code> , error 3703 occurs with state 10.
11	If the database being dropped has referential dependencies, Adaptive Server tries to drop the references from the other databases. If the ID of the reference from the other database is not found in <code>sysreferences</code> , error 3703 occurs with state 11.

Action	Recovering from error 3703 requires manually modifying one or more system tables. Call Sybase Technical Support for assistance in determining which tables to modify and how to do that.
Additional information	Before calling Sybase Technical Support, have the information available that is listed in "Reporting errors" on page 4.
Version in which this error is raised	All versions

This section contains error messages for Adaptive Server transaction processing.

Error	Page
Error 3904	445
Error 3905	447
Error 3906	448
Error 3908	450
Error 3917	451
Error 3949	453

Error 3904

Severity

21

Message text

Can't unsplit logical page %ld in object '%S_OBJID' in database '%S_DBID' - both pages together contain more data than will fit on one page.

Explanation

Data and index pages are split when Adaptive Server determines that a row cannot be inserted due to insufficient room on the existing page. During transaction rollback (either by a user request or during database recovery), Adaptive Server tries to undo the changes done by the transaction. Error 3904 occurs when Adaptive Server fails to unsplit pages that were split during the transaction. Error 3904 is caused by an Adaptive Server problem.

When this error occurs, data integrity may have been compromised.

- Action
- 1 Use the procedure in “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* to identify which table and index correspond to the page number in the error message text.

Warning! The procedure referenced in step 1 uses the `dbcc page` command. Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 If the object encountering the error is not a system table (a system table’s object ID is less than 100), go to step 3.

If the object with the error is a system table and the index ID is not 0 (zero), refer to “Fixing a corrupted system table index” in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide for instructions on how to repair the system table index.

If the object with the error is a system table and the index ID is 0, contact Sybase Technical Support to help you repair the corruption. However, it may be necessary to restore from clean backups.
- 3 For a user table, if the index is 0, go to step 4. If the user table index ID is not 0:
 - a Run the `sp_helpindex` procedure on the index to get information for rebuilding it.
 - b Drop the index.
 - c Re-create the index. This clears the corruption in most cases.
 - d Run `dbcc checktable` on the table to verify data integrity. If errors occur, call Sybase Technical Support.
- 4 For a user table, if the index ID is 0, run `dbcc checktable` on the table to verify data integrity. If errors occur, restore the database from clean backups.

Additional information Refer to “drop index” and “create index” in the *Reference Manual: Commands* for information about dropping and re-creating indexes.

Refer to “Developing a Backup and Recovery Plan” in the *System Administration Guide: Volume 2* for information about how to safely create, dump, load, and re-create databases

Version in which this error is raised

All versions

Error 3905

Severity

21

Message text

Can't unsplit logical page %ld in object '%S_OBJID' in database '%S_DBID' - row number %d is used on both pages.

Explanation

When Adaptive Server determines that a row cannot be inserted due to insufficient room on the existing page, data and index pages are split to accommodate the insertion. If the transaction is rolled back (either by a user request or during database recovery), Adaptive Server tries to undo the changes done by the transaction. The undo steps include restoring the offset table that contains the location of rows on the page.

Error 3905 occurs when Adaptive Server is unable to restore the offset table on a data page, therefore failing to unsplit pages that were split during the transaction.

Error 3905 is due to an Adaptive Server problem and is sometimes seen when buffer pools larger than 2K are in use. When this error occurs, data integrity may have been compromised.

Action

Use the steps described below to recreate the affected table, then call Sybase Technical Support for help in resolving the server problem:

- 1 If the object in question is a user table, go to step 2.

If the object is a system table (a system table name starts with “sys” and the object ID is less than 100), call Sybase Technical Support. They may be able to help you repair the corruption, but it may be necessary to restore from clean backups.

- 2 Use `sp_help (<table_name>)` to determine the table's lock scheme: all-pages-locked (APL) table or data-only-locked (DOL) table.
- 3 Identify the rows that correspond to the affected page:

```
dbcc page(<dbid>, <page_num>, 1)
```

Where *<dbid>* and *<page_num>* are the database ID and page number shown in the error message respectively.

- 4 Use one of the following methods to recover from the problem:
 - Select the table data into a new table, drop the old table, and rename the new table to the old table name. Omit the affected rows from step 3 when selecting data from the old table.

For a DOL table, if there is an index on the table, force this index when copying data.
 - Bulk copy the affected table out, drop and re-create the table, and bulk copy the table back in. This is the most efficient solution for a very large table. Omit the affected rows from step 3 when copying the table back in.
- 5 Run `dbcc checktable` to check the new table.

Additional information

Avoid using named caches or large I/O buffer pools until the problem is resolved.

When you call Sybase Technical Support, have the information ready that is listed in “Reporting errors” on page 4.

Version in which this error is raised

All versions

Error 3906

Severity

16

Message text

Attempt to BEGIN TRANsaction in database '%.*s' failed because database is READ ONLY.

Explanation

Error 3906 is raised when a transaction, explicit or implicit, is attempted while the database is in read-only mode.

Error 3906 is raised with the following states:

State	Meaning
1	Attempted to start a transaction when the database is in read-only mode. Applies to all possible transaction types.
2	Attempted to start a subordinate transaction (child transaction started by a parent transaction in a parallel setting) when the database is in read-only mode.

Error 3906 is raised in two distinct cases, described below.

Case 1 – User transaction attempted in read-only database

A user attempted to start a transaction while the database is in read-only mode. This is not allowed.

Case 2 – Re-resolution in a read-only database

The error may occur when a view or procedure must be resolved in a read-only database. Re-resolution must update the `sysprocedures` table, but since the database is read-only, error 3906 is raised. This scenario can occur during a load into a read-only database, when views and procedures must be re-resolved.

Action

Corrective action depends on the scenario in which error 3906 was raised, as described in the above cases.

Case 1 – User transaction in read-only database (all versions)

Re-try the transaction when the database is no longer read-only.

Case 2 – Object re-resolution in a read-only database (all versions)

If error 3906 was raised when you attempt to select from a view or execute a stored procedure, correct the problem by making the database temporarily available for writes and re-compiling objects. Keep in mind, however, that if the database is used as a warm standby, this activity will create additional log records, which will invalidate the database as a warm standby.

To re-resolve compiled objects in the database:

- 1 Put the database in read-write mode.

```
1> sp_dboption <database_name>, "read only", false
2> go
```

- 2 Execute all compiled objects such as stored procedures and views once.
- 3 Return the database to read-only mode.

```
1> sp_dboption <database_name>, "read only", true
2> go
```

Additional information

Starting with Adaptive Server Enterprise version 11.9.x, which introduced standby access mode to allow access to the database during load sequences, you can run `alter database` while in standby access. However since transactions cannot be started in this mode, you must re-establish the last chance threshold (LCT) after the database is fully online following the completion of the load sequence. Follow the steps in case 2 above.

Version in which this error is raised All versions

Error 3908

Severity 16

Message text Attempt to BEGIN TRANSACTION in database '%.*s' failed because database is in BYPASS RECOVERY mode.

Explanation Adaptive Server can start various types of transactions such as local (in response to a local command), internal (for the server's own use), external (started by an external client like a TP Monitor) and subordinate (child transactions started by a parent transaction in parallel).

Bypass recovery starts Adaptive Server without recovering one or more databases. A database in this mode can be accessed even though it is not recovered. Use this mode to allow access to the database for problem analysis, to copy out data, and so on.

Error 3908 is raised when a transaction is attempted while the database is in bypass recovery. Although certain commands are permissible in this mode, no transactions—regardless of type—can be started when the database is in bypass recovery.

Error 3908 is raised with the following states:

State	Meaning
1	Attempted to start a transaction when the database is not available for update. Applies to all possible transaction types.
2	Attempted to start a subordinate transaction when the database is not available for update.

Action Do not attempt to start a transaction while the database is in bypass recovery. Check if any automated jobs have submitted a batch job that attempted such a transaction. First restore the database to the original working status it had (prior to bypass status) before attempting any transactions.

Use the following steps to reset the database status:

Note Although a status of zero is usually appropriate, it is not the only possible status value you can use. Reset the database to the actual, normal status it had prior to entering bypass recovery status.

1 Reset the database status to 0:

```
1> sp_configure "allow updates", 1
2> go

1> begin tran
2> go

1> use master
2> go

1> update sysdatabases set status=0
2> where name=<database_name>
3> go
```

Verify that the above update command affected only one row (if more than one row was affected, issue a rollback transaction), then commit the transaction and shut down Adaptive Server:

```
1> commit tran
2> go

1> shutdown with nowait
2> go
```

2 Restart Adaptive Server. Run dbcc checkdb and dbcc checkalloc on the affected database to make sure there are no errors.

Version in which this error is raised

All versions

Error 3917

Severity

16

Message text

An attempt was made to startup a multi-database update in database '%S_DBID' after starting up the master transaction in tempdb. This could render the database

'%S_DBID' unrecoverable and this update is therefore not allowed.

Explanation Error 3917 is raised when you initiate a transaction in `tempdb` and the transaction attempts to update one or more tables in another database. For example:

```
1> use tempdb
2> go

1> begin tran
2> go

1> update geo..nation
2> set n_comment="This is a grade V earthquake zone"
3> where n_nationkey = 4
4> go
```

Msg 3917, Level 16, State 1:

Line 1:

An attempt was made to startup a multi-database update in database 'geo' after starting up the master transaction in `tempdb`. This could render the database 'geo' unrecoverable and this update is therefore not allowed.

You cannot start update transactions in `tempdb` because `tempdb` is truncated upon recovery. If the transaction is left in a prepare state and it becomes necessary to restart the server, this will render the user database (the one being updated) unrecoverable.

Error 3917 can also be raised when a `tempdb` transaction log gets filled up, and is accompanied by error 1105 (see “Error 1105” on page 173).

Action Check the Adaptive Server error log to help determine the circumstances in which the error is raised.

1 If error 3917 was raised when you attempted a transaction that contained an update or insert statement, the transaction was incorrectly initiated from `tempdb`. Take one of the following steps, as applicable:

- Add “`use <database_name>`” in the SQL script to change the database from `tempdb` to the database being updated, and retry the transaction;

or

- Change the default database for the user executing the transaction to the database where the insert is taking place.
- 2 If error 1105 in `tempdb` accompanies the 3917 error, an Adaptive Server problem probably occurred that prevents the checkpoint process from executing because the transaction log is filled up. Shut down and restart the server.

Additional information	When you call Sybase Technical Support, have the information ready that is listed in “Reporting errors” on page 4.
Version in which this error is raised	All versions

Error 3949

Severity	20
Message text	An illegal attempt was made to update `%. *s` in database `%S_DBID` using a transaction that was started in database `%S_DBID` (mode: %d, type: %d).
Explanation	This is an internal error and is raised only if trace flag 3945 is enabled. Any transaction that modifies an object will write log records in that database (associated with the object). When a transaction is started in one database and tries to modify an object in other database, this error is raised.
Action	It is not illegal to update an object in a different database than which the transaction was started; but Adaptive Server has detected an internal error when doing so. Disable trace flag 3945 to avoid seeing this message.
Additional information	If the error persists, a configurable shared memory dump (CSMD) may be helpful in determining the root cause. Information about CSMD can be found in the chapter “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide
Version in which this error is raised	Version 15.0 and later

This chapter contains general error messages for the Adaptive Server *dataserver* utility program.

Note The *dataserver* binary in Windows is *sqlsrvr.exe*. If you use a Windows platform, substitute all references to *dataserver* in this chapter with *sqlsrvr*.

Error	Page
Error 4001	455
Error 4002	456
Error 4020	459

Error 4001

Severity 11

Message text Cannot open default database '%.*s'.

Explanation Error 4001 is raised when Adaptive Server is unable to open the users default database at login.

After raising the error, the server verifies whether the user is allowed to log in to *master*. If the *master* database is forbidden for the user, the login request is denied.

Common reasons why the default database cannot be opened include:

- The database may not be available for use. For example, it may be in recovery.
- The database may be up, but not available for general use. For example, it may be undergoing failover as part of the High Availability (HA) subsystem.

	<ul style="list-style-type: none">• Permission checks failed. For example, if the database was created by loading from another server, but the login information in the source and target servers does not match, login may fail.
Action	<p>Retry the login in the event that the database was in recovery or temporarily unavailable.</p> <p>Since the 4001 error is usually a symptom of another problem, check the server error log for any other errors that may have been raised prior to 4001. Check for the following errors and retry:</p> <ul style="list-style-type: none">• 921 errors (Database '%.s' has not been recovered yet - please wait and try again.) indicate that the database is in recovery, or is marked suspect due to an earlier problem. Refer to the "Error 921" on page 155 for instructions.• 2206 errors (Database %.s with dbid %d is already shut down.) indicate that the database is not available for general use. These errors are raised in version 12.0 and later. Check for other HA subsystem messages to determine the cause.• 10351 errors (Server user id %d is not a valid user in database '%.s') indicate a permissions problem. If the database was loaded from another server, the problem may be due to an incorrect <code>sudo/uid</code> combination in <code>sysusers</code>. Check <code>syslogins</code> entries in the source and target server, respectively. Refer to "Step 12: Restore server user IDs" in the chapter "Restoring System Databases" the <i>System Administration Guide: Volume 2</i> for details.
Version in which this error is raised	All versions

Error 4002

Severity	14
Message text	Login failed.
Explanation	<p>This error occurs when a login request is rejected by Adaptive Server.</p> <p>Some reasons for error 4002 are:</p> <ul style="list-style-type: none">• The user name does not exist in <code>syslogins</code> (names are case-sensitive).• An incorrect password has been entered (passwords are case-sensitive).

- The account for the user trying to log in is locked.
- The user is trying to log in to an Adaptive Server:
 - That does not exist, or
 - To which the user does not have access, or
 - That is not running.
- Adaptive Server is in single-user mode and only “sa” can login.
- Adaptive Server shutdown is in progress and only “sa” can login.
- Remote Adaptive Server names are translated from the login record to a site ID. If the Adaptive Server name is not found in `sys.servers`, then the login fails.
- A remote login is being attempted and `sysremotelogins` is missing.
- A remote login is being attempted and the user cannot be found in `sysremotelogins`.
- For “sa,” your Replication Server password is different from your Adaptive Server password.

Action

Determine whether any of the reasons listed above are causing the error. If you have checked each of these and you still cannot log in to Adaptive Server, call Sybase Technical Support.

Following are instructions for checking on the above reasons:

- Login name does not exist in `syslogins` (names are case-sensitive).
 - When Adaptive Server is first installed, the only login is “sa” and it has no password.
 - If using `isql`, make sure the user is specifying the `-U` option and that the correct name and case are being specified.
 - Check `syslogins` to make sure the user is spelling the login name correctly:

```
1> select * from syslogins
2> go
```
- An incorrect password has been entered (passwords are case-sensitive).
 - If using `isql`, make sure the user is specifying the `-P` option and the correct password and case are being specified.
 - If a user cannot remember his or her password:

Change the password using `sp_password` (only the System Security Officer (SSO) can change someone else's password):

```
1> sp_password <caller_passwd>,  
2> <new_passwd>, <login_name>  
3> go
```

Where `<caller_passwd>` is the password of the login account that is currently executing `sp_password` and `<new_passwd>` is the new password for `<login_name>` (the name of the user whose password is being changed).

- If you cannot remember your “sa” password and you cannot get into Adaptive Server with any other login:
 - 1 Run `dataserver -p` to generate a new password for a System Security Officer’s account.
 - 2 Start Adaptive Server with `-p`.
 - 3 Immediately log in to Adaptive Server with the new random password, and execute `sp_password` to reset your password to a more secure one.

Refer to the Adaptive Server Enterprise core manuals for more information.

- The account for the user trying to log in is locked. To determine whether this is the case, type:

```
1> sp_displaylogin <login_name>  
2> go
```

Where `<login_name>` is the login name of the user. Check the value for the `status` column:

- If the value is 1, the password is less than 6 characters or NULL.
- If the value is 2, the account is locked.
- If the value is 3, the account is expired.

To unlock a user account:

```
1> sp_locklogin <login_name> "unlock"  
2> go
```

(Only System Administrators and System Security Officers can use `sp_locklogin` to lock and unlock logins.)

- The user is trying to log into a server that does not exist or to which the user does not have access. Check the user's environment file to make sure the value of *DSQUERY* is defined as the server to which the user is trying to log on.
- When Adaptive Server is in single-user mode, only "sa" can login.
 - Refer to "Returning Adaptive Server to multiuser mode" in the chapter "Server Recovery Tasks" in the most recent version of *Troubleshooting and Disaster Recovery*.
 - Check your *runserver* file to make sure the -m option is not specified.
- When Adaptive Server shutdown is in progress, only "sa" can login. Tell the user to wait until Adaptive Server has come back up.
- Remote server names are translated from the login record to a site ID. If the server name is not found in *syssservers*, then the login fails. Use *sp_helpserver* to list the servers in *syssservers*.
- If a remote login is being attempted and *sysremotelogins* is missing, then the login fails. To determine whether this is the case, type:


```
1> sp_helpdb sysremotelogins
2> go
```
- If a remote login is being attempted and the user cannot be found in *sysremotelogins*, then the login fails. To determine whether this is the case, type:


```
1> sp_helpremotelogin
2> go
```
- For "sa," your Replication Server password is different from your Adaptive Server password. Make sure the "sa" passwords are the same.

Additional information

For information about logins and the procedures mentioned in this write-up, refer to the *System Administration Guide: Volume 1* chapter "Introduction to Security" and the subsequent security chapters in that guide.

Version in which this error is raised

All versions

Error 4020

Severity

20

Message text	While reading characters from the client, the receiving buffer has been detected to be too small.
Explanation	<p>All client/server communication occurs over a network using packets. Clients initiate a connection to the server. The connection sends client requests and server responses.</p> <p>Error 4020 is raised when the server detects an overflow in the destination buffer while reading one or more characters from the client. The error breaks the connection to the server.</p> <p>Possible reasons for error 4020 include:</p> <ul style="list-style-type: none">• application errors• corrupt packets being transmitted• incorrect network configuration
Action	<ol style="list-style-type: none">1 Collect all relevant information about the circumstances in which the error occurred: character set differences on client and server, changed or new applications, recent changes in the network configuration, and so on.2 Check for coding errors in the application which may be causing incorrectly sized packets to be sent.3 Check for possible packet corruption within the network.4 Check the network configuration. For example, if there is 7-bit communication at one end, and 8-bit at the other end, check that the network is properly configured to handle this setup.5 Check that Adaptive Server is correctly configured to work with the underlying network packet size or buffer size on your system. Relevant server parameters include: <pre>default network packet size maximum network packet size total memory (pre-12.5) or max memory (version 12.5 and later)</pre>
Version in which this error is raised	All versions

View Resolution Manager Errors (4400s)

This chapter describes errors that result from working with SQL-derived tables.

Error	Page
Error 4430	461

Error 4430

Severity	19
Message text	The number of tables including the derived tables in the query exceeded the limit of %ld tables.
Explanation	The combined number of tables and derived tables (see “How SQL-derived tables work” below) in the query is too high.
Action	Reduce the number of tables, including derived tables, specified in the query that encounters the error.
Additional information	A SQL-derived table is defined by one or more tables through the evaluation of a query expression. A SQL derived table is used in the query expression in which it is defined and exists only for the duration of the query. It is not described in system catalogs or stored on disk.

How SQL-derived tables work

A SQL-derived table is created with a derived table expression consisting of a nested select statement, as in the following example, which returns a list of cities in the publishers table of the pubs2 database:

```
select city from (select city from publishers) cities
```

For example, the SQL-derived table from the subquery above is named cities and has one column titled city. The SQL-derived table is defined by the nested select statement and persists only for the duration of the query, which returns the following output:

```
city
-----
Boston
Washington
Berkeley
```

A more complete explanation of derived tables, and their syntax and use, can be found in the chapter “SQL Derived Tables” in the *Transact SQL User’s Guide*.

Version in which this error is raised

Version 12.5.1 and later

This section describes errors for the truncate table command.

Error	Page
Error 4716	463

Error 4716

Severity	16
Message text	Cannot truncate table '%.*s' because it is partitioned.
Explanation	<p>By default, Adaptive Server stores a heap table's data in one doubly linked chain of database pages. Adaptive Server inserts all new rows into the last page of the chain. A transaction holds an exclusive lock on the last page while inserting new rows, which can block other, concurrent transactions from being inserted into the table.</p> <p>The partition clause of the alter table command allows you to partition user tables that do not have a clustered index. Partitioning creates additional page chains on the table, each with its own last page. This reduces page contention for concurrent inserts, and can also reduce I/O contention if the table is distributed over multiple physical devices.</p> <p>Error 4716 is raised when you try to truncate a table that is partitioned.</p>
Action	<p>If you want to truncate the table, use the unpartition clause of the alter table command to concatenate all partitions:</p> <pre>1> alter table <table_name> unpartition 2> go</pre> <p>Then truncate the table.</p>
Additional information	Refer to the <i>Reference Manual: Commands</i> for information about the alter table command.
Version in which this error is raised	Versions earlier than 15.0

This section describes errors for the bulk copy utility (bcp). See the *Adaptive Server Enterprise Utility Guide* for more information about using the bcp utility.

Error	Page
Error 4801	465
Error 4806	466

Error 4801

Severity	20
Message text	Bulk_main: opentable on BULK INSERT table failed. Dbid=%d name='%.*s'.
Explanation	<p>Error 4801 is raised when you attempt to bulk copy data into a table, but bcp is unable to locate the target table.</p> <p>Possible causes include:</p> <ul style="list-style-type: none"> • A bcp syntax error • More than one session attempting to bcp into the same table • The target table was deleted or renamed by a concurrent user
Action	<ol style="list-style-type: none"> 1 Check the syntax of your bcp in command and the target table named in the command. 2 Verify that the target table exists, and that no other user is attempting to bulk copy to the table. 3 If you are making a substantial change, such as loading a large number of tables, you may need to increase the number of open objects configuration parameter. See the chapter “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for details.

Version in which this error is raised All versions

Error 4806

Severity 16

Message text You cannot run the non-logged version of bulk copy in this database. Please check with the DBO.

Explanation This error occurs when the `select/into bulkcopy` option is set to “false” and you use “fast” bulk copy into a table that has no indexes.

Note In newly created databases, the `select into/bulkcopy` option is set (by default) to the same as that in `model`.

Action Use the following options to recover from this error.

Turn on *select into/bulkcopy* (all versions)

Use the `sp_dboption` stored procedure to set the `select into/bulkcopy` option to “true”:

```
1> sp_dboption <database_name>,
2> "select into/bulkcopy", true
3> go

1> use <database_name>
2> go

1> checkpoint
2> go
```

Once a minimally logged operation such as “fast” bulk copy runs in the database, you cannot dump the transaction log to a device, because unlogged changes are not recoverable. Instead, perform a `dump database` as soon as possible to restore recoverability and allow transaction dumps to devices again.

Create an index for the table

Put an index on the table into which you are trying to copy. This causes inserts to be logged.

When you copy into a table that has indexes, a slower version of bcp is automatically used. The slow version, which does log data inserts in the transaction log, can cause the transaction log to become very large. You may need to use `dump transaction with truncate_only` until you can perform a full database dump. If you must do this, you will lose the ability to recover up-to-the-minute changes in data in case of a media failure. Also consider using smaller bcp batch sizes by adding the `-b` parameter to avoid a large size transaction log.

Version in which this error is raised

All versions

This section contains error messages for the `alter table` command.

Error	Page
Error 4951	469
Error 4953	470
Error 4954	471
Error 4956	471
Error 4964	472

Error 4951

Severity

16

Message text

Cannot create %d partitions on the table '%.*s' with clustered index because the maximum number of partitions allowed on a table with clustered index is %d.

Explanation

By default, Adaptive Server stores a heap table's data in one doubly linked chain of database pages. Adaptive Server inserts all new rows into the last page of the chain. A transaction holds an exclusive lock on the last page while inserting new rows, which can block other, concurrent transactions from being inserted into the table.

Partitioning creates additional page chains on the table, each with its own last page. This reduces page contention for concurrent inserts, and can also reduce I/O contention if the table is distributed over multiple physical devices.

The `partition` clause of the `alter table` command allows you to partition tables with or without a clustered index. However, after partitioning a clustered table, the entire page chain is placed in the first partition.

Error 4951 occurs when you attempt to partition a table with a clustered index, and exceed the maximum allowed partitions.

Action	Re-enter your partition command using a value for number of partitions that does not exceed the specified maximum.
Additional information	Refer to the <i>Reference Manual: Commands</i> for information about the alter table command.
Version in which this error is raised	All versions

Error 4953

Severity	16
Message text	Cannot partition table '%.*s' because it is already partitioned.
Explanation	<p>The partition clause of the alter table command allows you to partition user tables that do not have a clustered index. You cannot partition a table that is already partitioned.</p> <p>Error 4953 occurs when you try to partition a table that is already partitioned.</p>
Action	<p>To change the number of partitions for the table, follow these steps:</p> <ol style="list-style-type: none">1 Concatenate all existing page chains:<pre>1> use <database_name> 2> go 1> alter table <table_name> unpartition 2> go</pre>2 Repartition the table:<pre>1> alter table <table_name> partition 2> <number_of_partitions> 3> go</pre>
Additional information	Refer to the <i>Reference Manual: Commands</i> for information about the alter table command.
Version in which this error is raised	All versions

Error 4954

Severity	16
Message text	Cannot unpartition table '%.*s' because it is not partitioned.
Explanation	<p>The partition clause of the alter table command allows you to partition user tables that do not have a clustered index. The unpartition clause of the alter table command allows you to concatenate all existing page chains for these tables. You cannot unpartition a table that is not partitioned.</p> <p>Error 4954 occurs when you try to unpartition a table that is not partitioned.</p>
Action	<p>No action is required.</p> <p>If you want to determine whether a table is partitioned, use the following commands:</p> <pre> 1> use <database_name> 2> go 1> sp_help <table_name> 2> go </pre>
Additional information	Refer to the <i>Reference Manual: Commands</i> for information about the alter table command.
Version in which this error is raised	All versions

Error 4956

Severity	16
Message text	Cannot alter table if level 0 scans are active.
Explanation	<p>Isolation level 0 allows transactions to read uncommitted data (“dirty reads”). You can specify isolation level 0 for your queries using the at isolation syntax or you can specify level 0 for your session as part of the transaction isolation level option of the set command. Adaptive Server's default transaction isolation level is 1.</p> <p>Since alter table allows you to modify a table's schema, it would not be safe to allow the use of alter table when level 0 scans are active for that table. Error 4956 is raised if you enter an alter table command when the table in your command is currently being used for reading uncommitted data.</p>

Action	Try your alter table command again later, after the process that is using level 0 scans on the table is no longer active.
Additional information	Refer to the <i>Transact-SQL User's Guide</i> for information about selecting an isolation mode for either a query or for your session.
Version in which this error is raised	All versions

Error 4964

Severity	17
Message text	Lock scheme conversion failed due to insufficient SQL Server memory. Please retry later when there is less load/users on the SQL server, or ask your System Administrator to reconfigure SQL Server with more memory.
Explanation	<p>Adaptive Server provides two data locking schemes:</p> <ul style="list-style-type: none">• Datapages locking, which locks only the data pages• Datarows locking, which locks only the data rows <p>Since neither scheme locks index pages, they are referred to together as the data-only-locking (DOL) scheme.</p> <p>An additional locking scheme, known as all-pages-locking (APL), locks the data and index pages affected by queries. APL is the default locking scheme.</p> <p>When you use the <code>alter table</code> command to change a table's locking scheme from one lock scheme to other lock scheme, Adaptive Server creates memory structures to build a new table with the desired locking scheme and transfers data from the existing table.</p> <p>Error 4964 indicates that there is not enough memory available in Adaptive Server's data cache to effect the lock scheme change. The problem can occur at various stages of table conversion when:</p> <ul style="list-style-type: none">• Initializing a new object allocation map (OAM) page• Initializing a new index structure• Building the new table

Action	<p>Check the activity on your server and retry the lock scheme conversion when there is less load on the system.</p> <p>If the problem persists, check the memory allocation on your server. You may need to increase the memory allocated to the data cache.</p> <p>Use <code>sp_configure</code> to view the current values of memory-related parameters on your system:</p> <pre>1> sp_configure "Memory Use" 2> go</pre> <p>Refer to “Configuring Memory” in the <i>System Administration Guide: Volume 2</i> for details on monitoring and adjusting memory allocation parameters.</p>
Version in which this error is raised	All versions

This chapter contains error messages for the alter database command.

Error	Page
Error 5006	475
Error 5013	476
Error 5018	479
Error 5034	480

Error 5006

Severity	16
Message text	Could not find enough space on disks to extend database %.*s.
Explanation	<p>This error occurs when an alter database command refers to an Adaptive Server device that is completely full.</p> <p>If the device is not full but has less space than the alter database command requests, the database will be extended to fill up the available space on the device and no error message will be displayed. Instead, you will see the usual status message from alter database, but it will report fewer pages than you requested.</p> <p>This error message does not usually appear in the error log. However, if the exhausted device also uses up the available physical disk drive space as well as the space allocated for Adaptive Server, kernel messages declaring the file system full will occur in the error log. If this happens, fix the file system problem as well as the lack of space for Adaptive Server.</p>
Action	<p>To recover from error 5006, perform one of the following procedures:</p> <ul style="list-style-type: none"> Use a device name in the alter database command that is large enough. Use sp_helpdevice and sp_helpdb to find out which devices have enough room.

- Add a new device to Adaptive Server using the `disk init` command, then refer to the new device using the `alter database` command.
- Although there is no particular performance benefit to keeping a database entirely on one device, if you need to do this at a site, use this procedure:
 - a Back up all the databases on the device using the `dump database` command.
 - b Drop the device using `sp_dropdevice`.
 - c Initialize a larger device using `disk init`.
 - d Create the old databases on the new device using standard procedures. It is important to create the databases exactly like the old ones or errors will occur. Refer to “Error 2558” on page 314 for more information.
 - e Reload each database with the `load database` command.
 - f Use the `online database` command for each database to make the databases available for use.

Refer to “Developing a Backup and Recovery Plan” in the *System Administration Guide: Volume 2* for information about how to safely create, dump, load and re-create databases.

Version in which this error is raised

All versions

Error 5013

Severity	16
Message text	Cannot extend the MASTER database onto any device other than 'master'. The ALTER DATABASE was aborted.
Explanation	This error occurs when you try to extend the master database onto a device other than the master device.

Note It is recommended that you keep user objects out of the master database. If you keep user databases off the master device, you allow space in case the master database needs to grow. In addition, if you ever need to rebuild the master device, it will be easier if it does not contain user databases.

Action

Adaptive Server users can move any “home-grown” system procedures that start with “sp_” to sybssystemprocs (by dropping them from the master database and creating them in sybssystemprocs).

Extend the master database only if absolutely necessary! If you are sure you must increase the master database size and have no room on the current master device, use the following procedure to remove user databases from the master device.

Moving user databases

- 1 Dump the user databases with the `dump database` command.
- 2 Rename the dumped databases on the master device with `sp_renamedb`.
- 3 Re-create the databases with their original names on another device with `create database`. Be sure they are created exactly like the old databases, to avoid 2558 and other errors. Refer to “Error 2558” on page 314 for more information.
- 4 Load the dumps with `load database`.
- 5 Use the `online database` command for each database to make the databases available for use.
- 6 Check the databases in their new location to make sure the load was successful (that is, perform a simple query with `isql`), and if everything loaded successfully, drop the old databases from the master device.

You can now try to increase the size of the master database on the master device with the `alter database` command.

Increasing master device size

If the master device contains only the master database and the master device is too small, then use the following procedure:

Warning! Altering the master device is extremely risky! Avoid it if at all possible. Be familiar with the recovery methods in the chapter “Adaptive Server System Recovery” in the most recent version of *Troubleshooting and Disaster Recovery* in case you lose your master database or device.

- 1 Back up the master database with the `dump database` command.
- 2 Save the contents of key system tables such as `sysdatabases`, `sysdevices`, `sysusages`, and `syslogins`.

3 Use the `buildmaster` utility to build a new master device with enough extra space so that you will never need to increase the master device again. When `buildmaster` completes, a new master database will exist on the new master device.

4 Expand the size of the new master database with the `alter database` command, if necessary, so that it matches the size of the dumped master database.

5 Execute the following command in `isql`:

```
1> select name, high from master..sysdevices
2> where name = "master"
3> go
```

and note the “high” value for the master device.

6 Start Adaptive Server in single-user mode. Refer to “Starting Adaptive Server in single-user mode” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.

7 Load the dump of the master database. Test the database to make sure it is functioning correctly before deleting any extra backups of master.

8 Allow updates to the system catalog:

```
1> sp_configure, "allow updates", 1
2> go
```

9 Reset the “high” value in `master..sysdevices`:

```
1> begin transaction
2> go
1> update master..sysdevices
2> set high = <value_of_high_from_step_5>
3> where name = "master"
4> go
```

10 If the previous update affected only one row, commit the transaction.

11 Restart Adaptive Server.

12 Turn off allow updates:

```
1> sp_configure, "allow updates", 0
2> go
```

Version in which this
error is raised

All versions

Error 5018

Severity	17
Message text	Caution: You have set up this database to include space on disk %d for both data and the transaction log. This can make recovery impossible if that disk fails.
Explanation	<p>Error 5018 is raised when you execute <code>alter database</code> to add storage space, and the disk map for the database is found to contain both data and log segments.</p> <p>The message can also be raised in another situation. When you load a database from dump, Adaptive Server compares the segments on the dump with the device allocations in the target database. If they match, the server maintains the allocations. If, however, there is a mismatch, the load will remap segments to the free device allocations. The remapping may result in mixed data and log segments, raising the 5018 error.</p> <p>Why is recovery difficult when data and log are mixed? To understand this, consider first what happens when you keep data and log separate, and make a complete database backup (say) every 7 days. If the device containing the data crashes, you can restore the database from the latest backup, then apply the committed transactions from the log device, restoring the database to practically the last moments before failure.</p> <p>In contrast, if data and log are mixed and the device fails, you can only restore the database up to the time of the last full backup, since there is no way to apply the transaction log.</p>
Action	<p>If the error was raised when creating a database</p> <p>Unless you are creating small, non-critical databases, always place the log on a separate device. This allows you to dump transaction logs, ensures full recovery from hard disk crashes and has many other benefits including performance improvement. For details, refer to “Placing the Transaction Log on a Separate Device” in the chapter “Creating and Managing User Databases” in the <i>System Administration Guide: Volume 2</i>.</p> <p>if the error was raised when loading a database</p> <p>It is likely that there is a mismatch in the segment mappings between the dump’s source and target databases. To correct this problem:</p> <ol style="list-style-type: none"> 1 In the source server (where the dump is to be executed), run the following query: <pre> 1> use master 2> go 1> select * from sysusages where dbid=db_id("<database_name>") </pre>

2> go

Each line of output corresponds to a segment, with a *segmap* of 4 denoting a log segment, and any other *segmap* denoting a data segment. Save the output.

- 2 Dump the database.
- 3 Create the target database for load, specifying the segments in the same order and with the same sizes as defined by the query output in step 1. For simplicity, make sure that all non-log segments have a *segmap* value of 3; if you had user-defined segments, the load will automatically remap these segments.

Check that the segment mappings are the same, by running the following query for the new database:

```
1> select * from sysusages where dbid=db_id("<database_name>")
2> go
```

- 4 Load the database.

Version in which this error is raised

All versions

Error 5034

Severity

16

Message text

ALTER DATABASE was successful. Cannot modify the last-chance threshold because the database '%.*s' is in read-only mode.

Explanation

Each database that stores its transaction log on a separate segment has a last chance threshold (LCT). This is an estimate of the number of free log pages that would be required to back up the transaction log.

Error 5034 is raised when you use `alter database` to extend log space on a read-only database. The database is altered and log space added, but the LCT cannot be modified because this requires an update to `sys thresholds`, which is not possible in a read-only database.

Error 5034 is followed immediately by error 5035 ("Remember to execute `SELECT LCT_ADMIN('LASTCHANCE', %d)`, when the database is made writable.

Note Version 11.0.3.2 and earlier raise error 3906 ("Attempt to BEGIN TRANsaction in database '%.*s' failed because database is READ ONLY.") when `alter database` is used to extend log space. The database is altered and log space added, but the LCT cannot be modified. See "Error 3906" on page 448 for details.

Action

Re-establish the LCT after the database is available for updates using the following steps:

- 1 Use the following command to determine the new last-chance threshold:

```
1> use <database name>
2> go
```

```
1> select lct_admin("reserve", <log size in pages>)
2> go
```

For example, on a 20MB transaction log (containing 10240 2KB pages):

```
1> select lct_admin("reserve", 10240)
2> go
-----
616
```

The resulting value is the amount of space required to dump the transaction log.

- 2 Modify the LCT using:

```
1> select lct_admin('lastchance', <new threshold>)
2> go
```

Additional information

Refer to "Creating a Free-Space Threshold for the Log Segment" in the chapter "Managing Free Space with Thresholds" in the *System Administration Guide: Volume 2* for additional information on setting thresholds.

For databases normally operated in read-only mode, no updates occur on the database and so there is no need to update the LCT.

Version in which this error is raised

All versions

This section contains error messages that pertain to disk activation and system administrator disk commands.

Error	Page
Error 5115	483
Error 5123	485
Error 5142	487

Error 5115

Severity	16
Message text	I/O error during disk initialization. PHYSNAME '%.*s' may be incorrect or %ld (VSTART + SIZE) may exceed the size of the device. Please consult the SQL Server error log for more details.
<hr/>	
	Note This error may be caused by a hardware problem.
<hr/>	
Explanation	This error occurs when a disk init command fails. It can occur for the following reasons: <ul style="list-style-type: none"> • Incorrect physical device name • Incorrect size • Virtual device number not unique • Permission problem • UNIX kernel not configured for asynchronous I/O.
Action	<ol style="list-style-type: none"> 1 Check the Adaptive Server error log for information that might help determine the cause of the error. 2 If necessary, check the items mentioned in the sections below.

- 3 Make any needed changes and then run the command again.

Incorrect device name

Make sure you have the correct physical device name. Refer to the Adaptive Server installation and configuration guide for your platform or “Specifying a physical device name with disk init” in the *System Administration Guide: Volume 1* in the chapter “Initializing Database Devices” for instructions.

Incorrect size

The `size` parameter of the `disk init` command must be specified in units of 2K blocks in version 12.0.x and earlier; or Kilobytes (k or K), Megabytes (m or M), or Gigabytes (g or G) in version 12.5 and later.

Sizing 2K units

There are 512 2K blocks in 1MB. Use the formula:

$$\begin{aligned} \#_of_Megabytes * \#_of_data_pages_per_Megabyte = \\ \#_of_2K_blocks \end{aligned}$$

For example, to initialize a 10MB Sybase device, specify 5120 for the size:

$$10MB * 512 = 5120$$

To determine the number of 2K blocks to use for a raw partition on a UNIX system, convert the size of the partition, in sectors, to 2K blocks. The size of a sector varies on different UNIX systems. For example, it is 512 bytes on Sun, so use `sectors/4` for the `size` parameter. It is 1024 bytes on HP, so use `sectors/2`. Also, see the information below about virtual device numbers.

You must restart Adaptive Server or use a different virtual device number when you reissue the `disk init` command.

Refer to the Adaptive Server installation and configuration guide for your platform for more information about Sybase and your operating system’s partitions.

Note Not all operating systems have partitions as described above.

Virtual device number

The `vdevno` parameter of the `disk init` command must be unique. Also, if a `disk init` command fails, Adaptive Server saves the `vdevno` used, even though the `disk init` command failed. To make the `vdevno` available again, restart Adaptive Server. Use `sp_helpdevice` to determine which virtual device numbers have been used.

The available number of virtual devices is equal to the number of configured devices minus one. For example, if the number of devices configuration parameter is set to 10, then there are 9 available devices, 1 through 9.

Permission problem

If Adaptive Server was started using the “sybase” account, there must be permission to write to the raw partition or to create a file in the directory specified by `disk init`. Raw partitions used by Sybase should always be owned by the “sybase” account. Correct the permissions problem and execute the command again, using a different `vdevno`. Refer to “Virtual device number” on page 484 for more information.

UNIX kernel

If the UNIX kernel (on Sun Solaris 2.x systems only) is not configured for asynchronous I/O, the 5115 error may occur only if your master device is on a UNIX file, and you are trying to do a `disk init` on a raw partition. Have your UNIX administrator rebuild the kernel to allow for asynchronous I/O. Refer to “Error 823” on page 131 for more information.

Version in which this error is raised

All versions

Error 5123

Severity

16

Message text

```
DISK INIT encountered an error while attempting to
open/create the physical file. Please consult the SQL
Server error log (in the SQL Server boot directory) for
more details.
```

Note This error may be caused by a hardware problem.

Explanation

This error occurs when a `disk init` command fails while trying to initialize an operating system file or a raw partition for use as a database device. It can occur for the following reasons:

- Incorrect physical device name
- Permission problem

- Incorrect size specified (for raw devices, when you request initialization of a device larger than the size of the raw partition)
- Not enough room on the device.

Action

Check each section below, make any needed changes and then run the command again.

Incorrect device name

Make sure you have the correct physical device name. Refer to the Adaptive Server installation and configuration guide for your platform or “Specifying a physical device name with disk init” in the *System Administration Guide: Volume 1* in the chapter “Initializing Database Devices” for instructions.

Permission problem

If Adaptive Server was started using the “sybase” account, there must be permission to write to the raw partition or to create a file in the directory specified by `disk init`. Raw partitions used by Sybase should always be owned by the “sybase” account. Correct the permissions problem and execute the command again, using a different `vdevno`. Refer to “Virtual device number” for more information.

Note Adaptive Server should not be started by “root.” If it is started by root, you will get error 5123 if the device or file is owned by “sybase.”

Incorrect size (pre-12.5 servers)

When you request initialization of a raw device, you must specify a size that is not larger than the size of the raw partition.

The `size` parameter of the `disk init` command must be specified in units of 2K blocks (the size of most data pages). There are 512 2K blocks in a megabyte. Use the formula:

$$\#_of_Megabytes * \#_of_data_pages_per_Megabyte = \#_of_2K_blocks$$

For example, to initialize a 10MB Sybase device, specify 5120 for the size:

$$10MB * 512 = 5120$$

To determine the number of 2K blocks to use for a raw partition on a UNIX system, convert the size of the partition, in sectors, to 2K blocks. The size of a sector varies on different UNIX systems. For example, it is 512 bytes on Sun, so use sectors/4 for the size parameter. It is 1024 bytes on HP, so use sectors/2. Also, see the information below about virtual device numbers. You must restart Adaptive Server or use a different virtual device number when you reissue the disk init command.

Refer to the Adaptive Server Enterprise installation and configuration guide for your platform for more information about Sybase and your operating system's partitions. For example, on UNIX, there is a size limitation of 2GB for a device file.

Note Not all operating systems have partitions as described above.

Additional information

Refer to “Initializing Database Devices” in the *System Administration Guide: Volume 1* for information about disk init.

Version in which this error is raised

All versions

Error 5142

Severity

16

Message text

Mirroring for device '%.*s' is not currently enabled.

Explanation

Error 5142 is raised when you execute a disk unmirror command in retain mode, but the device in question is not currently mirrored.

Disk mirroring is automatically deactivated when a read or write to a mirrored device fails. You may see error 5142, possibly accompanied by additional messages at the console or in the error log if:

- There are transient device failures
- The mirror was set up improperly, for example if an incorrect device name was specified.

Action

Take the following steps:

- 1 Look at the server error log. Check for additional errors preceding the 5142 error to determine the probable cause. For example, an “udunmirror” kernel error indicates a physical disk problem leading to I/O failure; for corrective action, refer to the write-up for “udunmirror errors” on page 691 in this manual.
- 2 If the device was recently mirrored or remirrored, check the procedure used to set up the mirror. Since a database device and its mirror constitute one logical device, you should not initialize the mirror device with `disk init`. Also check the syntax used for the `disk mirror` command, ensuring that the path to the mirror is properly specified. Finally, check disk configuration and make sure the disk is properly labeled.

Additional information

To retain use of asynchronous I/O, always mirror devices that are capable of asynchronous I/O to other devices capable of asynchronous I/O. On platforms that do not support asynchronous I/O for file system devices, you must mirror raw devices to raw devices and file system files to file system files.

Refer to “Initializing Mirrors” in the chapter “Mirroring Database Devices” in the *System Administration Guide: Volume 2* for details.

Version in which this error is raised

All versions

Transaction Coordinator Errors (5600s)

This section contains write-ups for common Adaptive Server Enterprise Transaction Coordinator (ASTC) errors.

Error	Page
Error 5602	489

Error 5602

Severity	17
Message text	Could not find an available distributed transaction participant descriptor. Retry later, or please contact a user with System Administrator (SA) role to reconfigure your system for more descriptors.
Explanation	<p>A distributed transaction (DTX) participant is an internal memory structure that the Adaptive Server Transaction Coordination (ASTC) service uses to manage a remote transaction branch. As transactions propagate to remote servers, the ASTC service must obtain new DTX participants (also known as DTX descriptors) to manage those branches.</p> <p>Error 5602 is raised when there are not enough participants in the DTX descriptor pool to start or continue a remote distributed transaction. In-progress distributed transactions may abort due to a shortage of DTX participants.</p>
Action	<p>Update the number of dtx participants configuration parameter. To strike a balance between too many participants (wasting memory) and too few (failing transactions), use <code>sp_monitorconfig</code> during a peak period:</p> <pre>1> sp_monitorconfig "number of dtx participants" 2> go</pre>

If the *#Free* value is zero or very low, consider increasing the value to increase the number of remote transaction branches that Adaptive Server can handle.

Version in which this error is raised

Version 12.0 and later

This section contains error messages for Adaptive Server Enterprise Open Client applications.

Error	Page
Error 5701	491
Error 5704	492

Error 5701

Severity	10
Message text	Changed database context to '%.*s'.
Explanation	<p>Adaptive Server sends this informational message to a DB-Library or CT-Library client:</p> <ul style="list-style-type: none"> • When the client logs in (the message displays the default database) • When the client changes the current database <p>The message is displayed unless the client application suppresses the message via its message handler.</p>
Action	<p>To suppress this message, add the following line to your client's Adaptive Server message handler:</p> <pre>if (msgno==5701) return (0);</pre> <p>This message cannot be suppressed on the Adaptive Server side; it must be handled on the client side.</p>
Version in which this error is raised	All versions

Error 5704

Severity	10
Message text	Changed client character set setting to '%.*s'.
Explanation	<p>Adaptive Server sends this informational message to a DB-Library or CT-Library client:</p> <ul style="list-style-type: none">• When the client logs in (the message displays the default character set)• When the client changes the current character set

Action

If you want to suppress this message, add the following line to your client's Adaptive Server message handler:

```
if (msgno==5704) return (0);
```

Note In version 12.x and version 15.x clients, a flag in the login packet is used to suppress these messages being sent by the Adaptive Server.

Version in which this error is raised

All versions of Adaptive Server Enterprise using pre-12.x clients

Configuration Errors (5800s, 10800s)

This chapter contains error messages that related to Adaptive Server Enterprise configuration.

Error	Page
Error 5806	494
Error 5808	494
Error 5824	495
Error 5846	497
Error 5847	497
Error 5848	498
Error 5849	499
Error 5850	500
Error 5851	501
Error 5852	502
Error 5853	502
Error 5854	503
Error 5857	504
Error 5859	505
Error 5861	506
Error 5863	507
Error 5865	508
Error 5866	509
Error 5867	510
Error 5868	511
Error 5893	511
Error 10841	512
Error 10865	513

Error 5806

Severity	16
Message text	Warning: The file '%s' has been overwritten. To prevent this rename, move, or delete old configuration files. See System Administration Guide for details.
Explanation	When there is a change in the server configuration, Adaptive Server overwrites or updates the existing configuration information in <code><servername>.cfg</code> . Adaptive Server updates the configuration file by keeping a copy of the previous configuration information in <code><servername>.nnn</code> file, where “ <i>nnn</i> ” is the next highest available number. Adaptive Server scans the <code>\$SYBASE/\$SYBASE_ASE</code> directory for the current maximum suffix and creates the backup file <code><servername>.nnn</code> . If the current maximum suffix (in the directory) is greater than or equal to 999, then Adaptive Server restarts the numbering; specifically, <code><servername>.001</code> will be created by overwriting the existing file.
Action	If the existing backup copies of the configuration file (<code><servername>.nnn</code>) are unnecessary, they can be removed. Be careful not to remove the <code><servername>.cfg</code> , <code><servername>.bak</code> , and <code><servername>.krg</code> files that may be in that same directory. If the backup copy files are necessary, they should either be renamed, moved to a different directory, or backed up and then removed.
Version in which this error is raised	Version 11.0 and later

Error 5808

Severity	10
Message text	WARNING: Dynamic loading of caches and pools through loading a new file are not supported. However, the loadfile '%s' will be inspected for consistency. Refer to 'sp_cacheconfig' and 'sp_poolconfig' to create or alter pools and caches.
Explanation	Configuration options are either dynamic or static. Dynamic parameters go into effect as soon as you execute <code>sp_configure</code> . Static parameters require Adaptive Server to reallocate memory; therefore, they take effect only after Adaptive Server has been restarted.

Creating, dropping, and changing the size of data caches requires a restart of Adaptive Server for the configuration to take effect. Changes to data caches that take effect without a restart include changing the type, creating, dropping, and resizing memory pools with `sp_poolconfig`, changing the wash percentage of the pools, and binding and unbinding objects.

Error 5808, an informational message, is displayed while Adaptive Server is loading configuration information from a configuration file. Although dynamic loading of caches by using `sp_configure` to load a new configuration file is not supported, Adaptive Server checks cache configuration information as it is read from the file to make sure valid values are specified.

Action	<p>No action is required.</p> <p>To activate cache and pool entries already made in the configuration file, shut down and restart Adaptive Server.</p> <p>In the future, use <code>sp_cacheconfig</code> and <code>sp_poolconfig</code> to configure caches and pools and avoid the 5808 error. You must shut down and restart Adaptive Server to make cache changes take effect.</p>
Additional information	See <i>Reference Manual: Procedures</i> for information about <code>sp_cacheconfig</code> and <code>sp_poolconfig</code> .
Version in which this error is raised	All versions

Error 5824

Severity	16
Message text	Cannot reconfigure server to use sort order ID %d, because the row for its underlying character set (ID %d) does not exist in syscharsets.
Explanation	<p>The <code>syscharsets</code> system table contains one row for each character set and sort order defined for use by Adaptive Server. One of the sort orders is marked in <code>master..sysconfigures</code> as the default sort order, which is the only order actually in use.</p> <p>This error is raised when you are reconfiguring Adaptive Server to use a different character set and/or sort order, and one of the following is true:</p> <ul style="list-style-type: none"> • A corrupt index exists on <code>syscharsets</code> • The specified sort order or character set does not exist in <code>syscharsets</code>

The error occurs with the following states:

State	Meaning
1	Adaptive Server could not find the specified character set in <code>syscharsets</code> .
2	Adaptive Server could not find the specified sort order in <code>syscharsets</code> .

Action

- 1 Determine the character sets that exist in `syscharsets` (note that a character set is defined by `id`, where the sort order is defined by a combination of the sort order ID and the character set ID `csid`):

```
1> select name, id from master..syscharsets
2> where csid = 0
3> go
```

This query can return more than one row if alternate character sets are installed. In this case, check the `sysconfigures` table or use `sp_helpsort` to determine the default character set.

If error 5824 is raised with state 1, and the specified character set does not exist, you will need to install it. Refer to “Configuring Character Sets, Sort Orders, and Languages” in the *System Administration Guide: Volume 1*.

- 2 Using the character set from step 1, determine whether the specified sort order exists in `syscharsets`:

```
1> select name, id, csid, description
2> from master..syscharsets
3> where csid = [<charset_id>]
4> go
```

Where `<charset_id>` is the ID from step 1.

- 3 If the character set or sort order data exist in `syscharsets`, check the table integrity:

```
1> dbcc checktable (syscharsets)
2> go
```

If errors are returned, recover from them by using the sections in this manual that describe those errors.

- 4 If errors persist, contact Sybase Technical Support. Depending on the type of problem you found, you may be able to take corrective action, such as manually changing the character sets or rebuilding the index on `syscharsets`.

Additional information

Before calling Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the:

- Output of `dbcc checktable` and the contents of `syscharsets`

- Output of `charset`, if the `charset` utility was used to try and change the character set or sort order

Version in which this error is raised

All versions

Error 5846

Severity

16

Message text

Illegal value '%ld' specified for configuration option '%s'. The legal values are between '%ld' and '%ld'.

Explanation

Error 5846 occurs when you specify a value for a configuration parameter that is lower than the lowest value allowed or higher than the highest value allowed.

Note Adaptive Server checks to make sure you have not set configuration parameters to values outside the allowed ranges. However, if you set a parameter such as `number of open objects` to a value that is too high for the amount of memory you have configured and that value is within the allowed range, error 5846 will not be raised and you may not be able to start Adaptive Server. Use `sp_configure` with the `verify` or `read` option to point out this sort of problem after you make configuration parameter changes and before you restart Adaptive Server.

Action

Refer to “Setting Configuration Parameters” in the *System Administration Guide: Volume 1* for legal value ranges for configuration parameters. Reset the value of the parameter named in the 5846 error message.

Additional information

See *Reference Manual: Procedures* for information about `sp_configure`.

Version in which this error is raised

All versions

Error 5847

Severity

16

Message text

Configuration Error: Configuration file '%s' does not exist or you do not have the required read/write

	<p>permission for the file.</p>
Explanation	<p>You can set or change the value of Adaptive Server configuration parameters in one of two ways:</p> <ul style="list-style-type: none">• By executing the system procedure <code>sp_configure</code> with the appropriate parameters and values.• By directly editing your configuration file and then invoking <code>sp_configure</code> with the <code>configuration file</code> option. <p>Error 5847 occurs when Adaptive Server cannot find the configuration file you specified in an <code>sp_configure</code> command or the user who started Adaptive Server does not have the required operating system-level read or write permission for the file.</p>
Action	<p>Check the name of the configuration file you specified in the <code>sp_configure</code> command. If the name was wrong, enter the command again with a valid file name.</p> <p>If the file name was not the problem, check the permissions on the configuration file you specified in the <code>sp_configure</code> command. If the user who started Adaptive Server does not have the required operating system-level read or write permission for the file, ask your System Administrator to change the protection on the file. Or have a different user, who has the appropriate access to the configuration file, start Adaptive Server.</p>
Additional information	<p>Refer to “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for information about configuration parameters and files.</p> <p>Refer to the <i>Reference Manual: Procedures</i> for information about <code>sp_configure</code>.</p>
Version in which this error is raised	All versions

Error 5848

Severity	16
Message text	<pre>Warning: Unknown parameter '%s' found on line '%ld' of configuration file '%s'. This parameter is ignored. Edit the configuration file to remove or correct this parameter.</pre>

Explanation	<p>You can set or change the value of Adaptive Server configuration parameters in one of two ways:</p> <ul style="list-style-type: none"> • By executing the system procedure <code>sp_configure</code> with the appropriate parameters and values. • By directly editing your configuration file and then invoking <code>sp_configure</code> with the <code>configuration file</code> option. <p>Error 5848 occurs during startup or run time when Adaptive Server detects a line in the configuration file that contains an unknown parameter. Adaptive Server ignores the unknown parameter.</p>
Action	Edit your configuration file. Using the parameter name and line number specified in the error message, remove or correct the incorrect entry.
Additional information	Refer to “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for information about configuration parameters and files.
Version in which this error is raised	All versions

Error 5849

Severity	16
Message text	Verification failed for parameter '%s'.
Explanation	<p>Adaptive Server runs a verification check on configuration parameters in the following circumstances:</p> <ul style="list-style-type: none"> • When you use <code>sp_configure</code> with the <code>verify</code> option to perform validation checking on the values in the configuration file you specify. This can help you avoid configuring Adaptive Server with invalid configuration values. • When you use <code>sp_configure</code> with the <code>read</code> option to perform validation checking on the values in the configuration file you specify and to read those values that pass validation into Adaptive Server. If any parameters are missing from the configuration file you specified, current values for those parameters are used. <p>Error 5849 occurs in situations when Adaptive Server detects a verification error for the parameter named in the error message.</p>

Action	Use <code>sp_configure</code> or edit your configuration file to correct the incorrect value for the parameter named in the error message. Run the verification check again to confirm that the problem has been corrected.
Additional information	Refer to “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for information about configuration parameters and files.
Version in which this error is raised	All versions

Error 5850

Severity	16
Message text	Notification failed for parameter '%s'.
Explanation	<p>You can use <code>sp_configure</code> with the <code>verify</code> or <code>read</code> option to perform validation checking on the parameters in a configuration file.</p> <p>Configuration parameters are either dynamic or static. Dynamic parameters go into effect as soon as you execute <code>sp_configure</code>. Static parameters require Adaptive Server to reallocate memory, and thus they take effect only after Adaptive Server has been restarted.</p> <p>During validation checking, if the value of a dynamic parameter is different in the configuration file than the current run value, Adaptive Server notifies itself that it should perform validation checking on the new value and notifies the user by writing an entry in the error log. Error 5850 occurs when that notification fails. It is caused by an Adaptive Server problem.</p>
Action	Try the <code>sp_configure</code> with <code>verify</code> or <code>read</code> command again. If error 5850 occurs again, call Sybase Technical Support.
Additional information	<p>Refer to “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for information about configuration parameters and files.</p> <p>Before calling Technical Support, have the information available that is listed in “Reporting errors” on page 4, including a copy of the configuration file being used.</p>
Version in which this error is raised	All versions

Error 5851

Severity	16
Message text	Configuration Error: Configuration file '%s' has an unknown format on line %d.
Explanation	<p>You can set or change the value of Adaptive Server configuration parameters in one of two ways:</p> <ul style="list-style-type: none"> • By executing the system procedure <code>sp_configure</code> with the appropriate parameters and values. • By directly editing your configuration file and then invoking <code>sp_configure</code> with the <code>configuration file</code> option. • By starting Adaptive Server with the <code>-c</code> option to specify a configuration file. <hr/> <p>Note For some parameters in the configuration file, the keyword “default” must be in capital letters (for example, for <code>wash size</code>).</p> <hr/> <p>The syntax for parameters in the configuration file is:</p> <pre><parameter_name>={<value> DEFAULT}</pre> <p>where:</p> <ul style="list-style-type: none"> • <i>parameter_name</i> – the name of the parameter you want to specify. • <i>value</i> – the numeric value to which you want to set <i>parameter_name</i>. • <code>DEFAULT</code> – specifies that you want to use the default value for <i>parameter_name</i>. If you specify <code>DEFAULT</code>, do not include <i>value</i>. <p>Error 5851 occurs during start-up or when you run <code>sp_configure</code> with the <code>verify</code> or <code>read</code> option when Adaptive Server detects a syntax error in a line in the configuration file. Adaptive Server ignores the line with the syntax error.</p>
Action	Edit your configuration file. Using the configuration file name and line number specified in the error message, remove or correct the incorrect entry.
Additional information	Refer to “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for information about configuration parameters and files.
Version in which this error is raised	All versions

Error 5852

Severity	16
Message text	Changing the value of '%s' is not allowed since it is a static option.
Explanation	<p>You can set or change the value of Adaptive Server configuration parameters in one of two ways:</p> <ul style="list-style-type: none">• By executing the system procedure <code>sp_configure</code> with the appropriate parameters and values.• By directly editing your configuration file and then invoking <code>sp_configure</code> with the <code>configuration file</code> option.• By starting Adaptive Server with the <code>-c</code> option to specify a configuration file. <p>Configuration parameters are either dynamic or static. Dynamic parameters go into effect as soon as you execute <code>sp_configure</code>. Static parameters, on the other hand, take effect only after Adaptive Server has been restarted. In version 12.5, several memory-related parameters were converted from static to dynamic.</p> <hr/> <p>Note Not all parameters listed here may be available in your Adaptive Server version.</p> <hr/> <p>Error 5852 occurs when you run <code>sp_configure</code> with the <code>read</code> or <code>verify</code> option and Adaptive Server determines that you have changed the value of the static option mentioned in the error message and have not yet shut down and restarted Adaptive Server.</p>
Action	No action is required unless you want the value of the static parameter to be used now. If you do, shut down and restart Adaptive Server.
Additional information	Refer to “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for information about configuration parameters and options.
Version in which this error is raised	All versions

Error 5853

Severity	16
----------	----

Message text	'%s' has been modified from '%ld' to '%ld' by the verification routine.
Explanation	<p>Adaptive Server runs a verification check on configuration parameters in the following circumstances:</p> <ul style="list-style-type: none"> • When you use <code>sp_configure</code> with the <code>verify</code> option to perform validation checking on the values in the configuration file you specify. This can help you avoid configuring Adaptive Server with invalid configuration values. • When you use <code>sp_configure</code> with the <code>read</code> option to perform validation checking on the values in the configuration file you specify and to read those values that pass validation into Adaptive Server. If any parameters are missing from the configuration file you specified, current values for those parameters are used. <p>Error 5853 is an informational message that is displayed in one of those situations when Adaptive Server changes the value for a configuration parameter because of rounding of the value. Rounding can occur when Adaptive Server needs to align to page boundaries.</p>
Action	No action is required.
Additional information	Refer to “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for information about configuration parameters and files.
Version in which this error is raised	All versions

Error 5854

Severity	18
Message text	Configuration Error: Cannot write out file '%s' due to system error '%s'.
Explanation	<p>You can set or change the value of Adaptive Server configuration parameters in one of two ways:</p> <ul style="list-style-type: none"> • By executing the system procedure <code>sp_configure</code> with the appropriate parameters and values. • By directly editing your configuration file and then invoking <code>sp_configure</code> with the <code>configuration file</code> option.

- By starting Adaptive Server with the -c option to specify a configuration file.

Adaptive Server:

- Writes information to the configuration file when you change a parameter using `sp_configure` or other stored procedures such as `sp_cacheconfig`.
- Writes information to the configuration file when you run `sp_configure` with the `verify` or `read` option and Adaptive Server needs to change the value for a configuration parameter because of rounding of the value.
- Creates a new configuration file when you use `sp_configure` with the `write` and `restore` options.

Error 5854 occurs when Adaptive Server is unable to write information to the configuration file in the error message because of the system error in the error message. The system error is an error returned by the operating system such as permission denied or no such directory or file.

Action	Refer to your operating system documentation for information about resolving the system error. Resolve the problem that caused the system error and try your action again.
Additional information	Refer to “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for information about configuration commands and files.
Version in which this error is raised	All versions

Error 5857

Severity	0
Message text	Changing the value of '%s' is not allowed since its value is calculated or is read only.
Explanation	Adaptive Server calculates the value for a number of configuration parameters. Some other configuration parameters are read only. Calculated parameters include: <ul style="list-style-type: none"> • allow syb_sendmsg function • executable codesize + overhead • syb_sendmsg port number

- total data cache size

Read-only parameters include:

- o/s asynch i/o enabled
- o/s file descriptors
- master device name
- master device mirror name
- configuration file

Error 5857 occurs when you try to change the value of the parameter named in the error message, since that parameter is either calculated by Adaptive Server or it is read-only.

Action	No action is necessary.
Additional information	Refer to “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for information about configuration parameters.
Version in which this error is raised	All versions

Error 5859

Severity	16
Message text	The configuration number '%d' is invalid.
Explanation	<p>Each Adaptive Server configuration parameter has a unique identification number assigned to it. For parameters set by users, this number is stored in the config column of the sysconfigures table.</p> <p>Adaptive Server uses these configuration numbers to identify the option to be displayed, loaded, modified, or verified when you run <code>sp_configure</code>. Error 5859 occurs when the configuration number Adaptive Server tries to use to identify the option is invalid.</p> <p>This error can be caused by corruption of <code>sysconfigures</code> or by an Adaptive Server problem.</p>
Action	When you call Sybase Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the output from <code>select * from sysconfigures</code> .

Version in which this error is raised All versions

Error 5861

Severity 16

Message text **Version 12.5 and later**

The current 'max memory' value '%ld', is not sufficient to change the parameter '%s' to '%ld'. 'max memory' should be greater than 'total logical memory' '%ld' required for the configuration.

Version 12.0.x and earlier

The 'total memory', '%ld', is not enough to change the parameter '%s' to '%ld'.

Explanation In versions 12.0.x and earlier, the total memory configuration parameter sets the size of memory, in 2K units, that Adaptive Server allocates from the operating system.

Error 5861 occurs when you try to change the value of the parameter named in the error message to the value shown in the error message because the value of the total memory parameter is not high enough to support the change.

If the error occurs when you are running `sp_configure`, the value you asked to be changed is not changed. If the error occurs during start-up (after you modified the configuration file), Adaptive Server does not start.

Action In version 12.0.x and earlier, if the error occurred when you were running `sp_configure`, specify a lower value for the parameter you are changing or raise the value of the total memory parameter. Then enter your original command again.

If the error occurred while you were starting Adaptive Server, edit the configuration file to make the changes, and then start Adaptive Server.

In version 12.5 and later, check that the memory available to the server is large enough to accommodate the total logical memory required by your configuration.

Additional information Refer to "Setting Configuration Parameters" in the *System Administration Guide: Volume 1* for information about configuration parameters.

See “Configuring Memory” in the *System Administration Guide: Volume 2* to help you determine appropriate values for memory-related configuration parameters.

Version in which this error is raised

All versions

Error 5863

Severity 18

Message text Invalid group encountered.

Explanation Adaptive Server Enterprise allows you to use `sp_configure` with a group name to display values for the configuration parameters belonging to that group. In version 12.5 and later, groups are:

- Backup/Recovery
- Cache Manager
- Component Integration Services administration
- DTM administration
- Diagnostics
- Disk I/O
- Error log
- Extended stored procedures
- General information
- Java services
- Languages
- Lock Manager
- Memory use
- Meta-data caches
- Monitoring
- Network Communication
- O/S Resources

- Physical memory
- Physical resources
- Processors
- Query tuning
- RepAgent thread administration
- SQL Server administration
- Security related
- Unicode
- User Environment

When Adaptive Server retrieves information for a configuration parameter group, it attempts to verify the group name. Error 5863 occurs when Adaptive Server is unable to verify the group name.

This error is caused by an Adaptive Server problem.

Action Call Sybase Technical Support with the information listed in “Reporting errors” on page 4 on hand.

Version in which this error is raised All versions

Error 5865

Severity 18

Message text `error in os_get_cpu_count: %d`

Explanation Before Adaptive Server checks the configuration parameters `min online engines` and `current number online engines` to make sure they contain valid values, it checks the number of online CPUs on the machine. Error 5865 occurs when Adaptive Server queries the operating system for the number of online CPUs and the operating system returns an error.

This error is caused by an operating system problem.

Action If the error occurred while starting Adaptive Server, restart Adaptive Server. Retry the command you were running when the error occurred.

If you are still getting the 5865 error, call Sybase Technical Support. Have the information listed in “Reporting errors” on page 4 on hand.

Version in which this error is raised

All versions

Error 5866

Severity

10

Message text

Configuration file '%s' has been written and the previous version has been renamed to '%s'.

Explanation

Adaptive Server:

- Writes information to the configuration file when you change a parameter using `sp_configure`.
- Writes information to the configuration file when you run `sp_configure` with the `verify` or `read` option and Adaptive Server needs to change the value for a configuration parameter because of rounding of the value.
- Creates a new configuration file when you use `sp_configure` with the `write` and `restore` options.

Each time you modify a configuration parameter with `sp_configure`, a new configuration file is created, using the naming convention `file_name.001`, `file_name.002`, `file_name.003` and so on, up to `file_name.999`.

Error 5866 is an informational message that is written to the Adaptive Server error log when Adaptive Server writes information to the configuration file in the message and renames the previous version of the file.

Action

No action is required.

Additional information

Refer to “Setting Configuration Parameters” in the *System Administration Guide: Volume 1* for information about configuration commands and files.

Version in which this error is raised

All versions

Error 5867

Severity 16

Message text The wash size for the %dk buffer pool in cache %s has been incorrectly configured. It must be a minimum of %d buffers and a maximum of %d percent of the number of buffers in the pool.

Explanation Wash size is the point in the cache at which Adaptive Server writes dirty pages to disk. You can use `sp_poolconfig` to place an entry in the configuration file to specify a wash size for a memory pool.

The default value for wash size depends on your server version. For example, in 11.5 the default wash size for a pool size less than 300MB is 20 percent of the buffers in the pool. (A buffer is a block of pages equal to the I/O size of the pool. All pages in a buffer are read from disk, written to disk, or flushed from the cache simultaneously.)

The minimum value for wash size is 10 buffers and the maximum value is 80 percent of the size of the pool.

Error 5867 occurs when you specify a value for wash size that is smaller than the minimum allowed size or larger than the maximum allowed size.

This error occurs with the following states:

State	Meaning
1	Occurs when you run <code>sp_configure</code> with the <code>verify</code> or <code>read</code> option after specifying a value for wash size that is smaller than the minimum allowed size or larger than the maximum allowed size.
2	Occurs when you use <code>sp_poolconfig</code> or update the configuration file directly to specify a value for wash size and that value is smaller than the minimum allowed size or larger than the maximum allowed size.

Action Determine a valid value for wash size for the buffer pool and cache specified in the error message and run your command again.

Additional information Refer to the *Reference Manual: Procedures* for information about `sp_poolconfig` and wash sizes.

Refer to “Configuring Data Caches” in the *System Administration Guide: Volume 2* for information about the wash area and wash configuration.

Version in which this error is raised All versions

Error 5868

Severity	16
Message text	Invalid type configured for cache %s. Cache type may be one of mixed OR log only but NOT both.
Explanation	<p>You can create, configure, reconfigure, and drop data caches by using the Adaptive Server procedure <code>sp_cacheconfig</code> or by modifying your configuration file.</p> <p>There are two types allowed for user-defined caches:</p> <ul style="list-style-type: none"> • <code>logonly</code> – only transaction log (syslogs) information is stored here. • <code>mixed</code> – both transaction log (syslogs) information and user data is stored here. This is the default. <p>Error 5868 occurs when Adaptive Server determines that you have tried to mix different types of cache. For example, error 5868 would occur if you try to define a cache as <code>log only</code> as well as <code>mixed</code> at the same time.</p>
Action	Use <code>sp_cacheconfig</code> or modify your configuration file to define a valid type for the cache named in the error message.
Additional information	Refer to the <i>Reference Manual: Procedures</i> for information about <code>sp_cacheconfig</code> and cache types.
Version in which this error is raised	All versions

Error 5893

Severity	16
Message text	The sum, (%s) + (%s %s) + (%s) + (%s), must be no greater than '%ld'.
Explanation	<p>Error 5893 is raised when the number of file descriptors available to Adaptive Server is less than the number needed by the process.</p> <p>When you start up Adaptive Server or run <code>sp_configure</code> to change certain configuration parameters (such as number of user connections), the following criterion must be met:</p> <pre>number of remote sites + max number network listeners + number of user connections</pre>

Error 10841

```
+ (number of devices * max online engines * 2)
<= o/s file descriptors - <descriptors for environment files>
```

Note On platforms that use light weight engines (Windows systems), use 1 instead of max online engines for the calculation.

The variable *<descriptors for environment files>* is not a configuration parameter. It accounts for file descriptors used by the server error log and other operational files. The *<descriptors for environment files>* value is 8 on Windows platforms and 10 on all other platforms.

Action	<p>Increase the available file descriptors or reduce the value of one or more configuration parameters in the above expression.</p> <p>Many operating systems allow you to configure the amount of file descriptors available per process. See the Adaptive Server installation and configuration guide for your platform for details on how to increase the available file descriptors.</p>
Additional information	<p>The <i>o/s file descriptors</i> parameter indicates the maximum per-process number of file descriptors configured for your operating system. This parameter is read-only and cannot be configured through Adaptive Server.</p>
Version in which this error is raised	<p>All versions</p>

Error 10841

Severity	<p>16</p>
Message text	<p>The value of parameter 'max memory' '%ld' cannot be less than the 'total logical memory' size '%ld'. Please reconfigure 'max memory' to be greater than or equal to the total logical memory required.</p>
Explanation	<p>The max memory configuration parameter lets you establish a maximum setting for the amount of memory that can be allocated to Adaptive Server. The total logical memory configuration parameter is the sum of the memory required for all the <i>sp_configure</i> parameters; while it is required to be available, the total logical memory may or may not be in use at any given moment.</p>

When Adaptive Server starts, it allocates memory needed for all nonconfigurable data structures and configurable parameters, including data and procedure caches. If you increase the server's memory configuration, for example by adding a worker process pool, the amount of memory available to configure must be able to absorb the added pool. Put another way, the new total logical memory (incremented by the addition of the pool) must not exceed max memory. If it does, error 10841 is raised.

Action	Reconfigure max memory to meet the new configuration requirements, or adjust other memory parameters so that they do not exceed the max memory limitation.
Additional information	Refer to "Setting Configuration Parameters" in the <i>System Administration Guide: Volume 1</i> for more information.
Version in which this error is raised	Version 12.5 and later

Error 10865

Severity	17
Message text	Unable to add memory to the cache (%!) as it is too fragmented. Changes will take effect when the server is rebooted. To add memory in the same instance delete the cache and recreate the cache with the required memory.
Explanation	This is a cache manager maintenance error related to fragmentation in the cache. This error is raised when the server is not able allocate the requested number of pages and buffers (requested by the DBMS tasks) or the size of the cache is increased more than 10 times.
Action	To remove the cache fragmentation, perform either of the following actions: <ul style="list-style-type: none"> • Delete the cache and reallocate it with the desired cache size; or • Reboot the Adaptive Server, at which point the cache allocation will increase to the new size.
Version in which this error is raised	Version 12.5.1 and later

This section contains error messages that are raised when the `kill` command is used to kill Adaptive Server processes.

Error	Page
Error 6103	515
Error 6107	516

Error 6103

Severity	17
Message text	Unable to do cleanup for the killed process; received Msg %d.
Explanation	<p>Upon termination of any Adaptive Server process, the server attempts to clean up by releasing resources used by the terminated process. Error 6103 occurs when the server is unable to complete this cleanup. The “<i>Msg %d</i>” appearing in the 6103 message text identifies the error that was received during cleanup.</p> <p>Error 6103 usually occurs after some other error or when an Adaptive Server process is aborted. This error can be serious, causing data corruption or an abnormal Adaptive Server shutdown.</p> <p>Error 6103 is usually caused by the errors that occurred before it or by an Adaptive Server problem.</p>
Action	<p>Review the error log to identify errors that occurred before the 6103 error. Use the information in this manual to recover from those errors. If the other errors are not covered in this manual, call Sybase Technical Support for assistance.</p> <p>To recover from error 6103 after reviewing other errors in the error log, shut down and restart Adaptive Server.</p> <p>If you cannot perform an Adaptive Server shutdown, kill the Server process at the operating system level.</p>

Run `dbcc checkdb` for the database that was being accessed by the killed process to determine whether data corruption has occurred. If other errors are displayed when `dbcc checkdb` is run, refer to the write-ups in this manual for those errors to determine how to recover from them.

If the 6103 error occurs again or resources (such as locks) are not released, call Sybase Technical Support.

Additional information

Refer to “kill” and “shutdown” in the *Reference Manual: Commands*. If you need to contact Sybase Technical support, have the information available that is listed in “Reporting errors” on page 4.

Version in which this error is raised

All versions

Error 6107

Severity

14

Message text

Only User processes can be KILLed or SYB_TERMINATED.

Explanation

A process is a task that is being carried out by Adaptive Server. Processes can be initiated by a user giving a command, or by Adaptive Server itself. You can see information about processes by running the system procedure `sp_who`.

Error 6107 is raised following an improper attempt to kill an Adaptive Server process; for example, it may be raised if you try to kill a background process. The error is raised with the following states:

State	Meaning
1	Attempted to kill a process that is not a user process (versions earlier than 12.0).
2	Attempted to kill a process that is not a user process or an agent process for quiesce database (version 12.0 and later).
3	Attempted to kill an agent process for quiesce database, but the <code>spid</code> requesting the kill is not the parent process of the agent (version 12.0 and later).

Action

Use `sp_who` to check the currently running processes and identify the process in question. Refer to “Killing Processes” in the *System Administration Guide: Volume 1* for a description of process status values and the effect of the `kill` command. If the error was raised when you attempted to clear a “killable” user process, take the following steps:

- 1 Perform a checkpoint in each database:


```
1> use <database_name>
```

```
2> go
```

```
1> checkpoint
```

```
2> go
```

2 Shut down the server using the `nowait` option:

```
1> shutdown with nowait
```

```
2> go
```

3 Restart the server.

Additional information

If you need to terminate a batch program or other application which logs in to Adaptive Server, always terminate the server process first using the `kill` command before attempting to terminate the batch at the operating system level.

Version in which this error is raised

All versions

Timestamp Errors (6900s)

This section contains error messages pertaining to Adaptive Server Enterprise timestamps.

Error	Page
Error 6901	519
Error 6902	520

Error 6901

Severity

21

Message text

Overflow on High component of timestamp occurred in database %d. Database table possibly corrupt.

Note This error may be caused by a hardware problem.

Explanation

Each database has one global timestamp that is kept in memory in a structure called a `dbtable`. The timestamp is not a date and time; rather, it is a sequence number (6 bytes, consisting of a high and low portion) which keeps track of modifications of pages within the database. Each time a page in the database is modified, the page gets the current timestamp and the global timestamp gets incremented. The global timestamp “travels” with a database through dumps, loads, and upgrades. Timestamps are critical during recovery, to ensure that transactions are properly rolled forward or rolled back. The global timestamp is also used for timestamp columns.

The maximum value for timestamps is very high (0xffff 0xffffffff), so it is very unlikely that under normal circumstances a database will reach this limit. However, use of some unsupported database maintenance commands can cause the global timestamp to reach the maximum value.

Each time `dbtable` is created for the database, Adaptive Server checks the timestamp value. Error 6901 is raised if the timestamp value exceeds the maximum.

Action

A database dump and load is not an option to address this problem since that does not reset the timestamp value. The only way to recover from error 6901 is to drop and re-create the database, using these steps:

- 1 `bcp` out all table data.
- 2 `defncopy` out all procedures, rules, defaults and triggers.
- 3 Use the `drop database` command to drop the database.
- 4 Use the `create database` command to re-create the database.
- 5 Create all tables.
- 6 `bcp` in all table data.
- 7 Create indexes.
- 8 `defncopy` in all procedures, rules, defaults and triggers.
- 9 Check that the user IDs in `sysusers` in the restored database are consistent with `master..syslogins`. Mismatched user IDs in databases can lead to permission problems.

Note Scripts are recommended for reliable re-creation of the database objects.

Additional information

Refer to “Developing a Backup and Recovery Plan” in the *System Administration Guide: Volume 2* for complete information about how to safely re-create databases.

Version in which this error is raised

All versions

Error 6902

Severity

21

Message text

Page timestamp value falls between the old and new timestamps from log. Page #=%ld, object id = %ld, page timestamp=%04x %08lx. Log: old timestamp=%04x %08lx,

```
new timestamp=%04x %08lx.
```

Note This error may be caused by a hardware problem.

Explanation

The objective of recovery algorithms in Adaptive Server is to bring all the databases to a consistent state: that is, each transaction is either completely applied to the database or no part of it is applied. (Database recovery occurs at Adaptive Server startup as well as when a load database or load transaction command is processed.)

Whenever a page is modified during run time, the page timestamp before the change and the page timestamp after the change are recorded in the log record for that modification. By comparing the timestamp on the page at recovery time with that of the timestamps in the log record, Adaptive Server can tell whether the page is in the state before or after the change was made to it.

Error 6902 occurs during Adaptive Server recovery if the timestamp of the page which is being recovered is neither the old log timestamp nor the new log timestamp. This is a fatal error and causes the recovery process to fail.

Error 6902 is probably a result of a more serious underlying error. Common causes of this error are:

- Data corruption during normal processing
- Abnormal shutdown of Adaptive Server during a transaction update
- Hardware errors

Action

Recover your database from backups or call Sybase Technical Support.

To prevent this error from occurring in the future, implement one of the following sections, depending on what caused the 6902 error.

Shutdown method

shutdown with `nowait` or abnormal shutdown of Adaptive Server should be avoided, if possible, after a heavy update period in your databases:

- Use `shutdown` to shut down Adaptive Server instead of `shutdown with nowait` after a period of heavy update activity in your databases.

or:

- `checkpoint` each database that is being used before shutting down Adaptive Server.

Hardware errors

Check the Adaptive Server error log to determine whether there are other indications of hardware problems, such as kernel messages reporting I/O errors.

Check the operating system error log or diagnostics utilities for I/O errors.

Using *dump transaction with no_log*

Using `dump transaction with no_log` can, in some circumstances, result in a 6902 error, especially when used after a heavy update period. Therefore, do not use `dump transaction with no_log` unless your transaction log is 100 percent full. Try using `dump transaction with truncate_only` first and use the `no_log` option only as a last resort.

If you find that the `no_log` option is used periodically, you might not be performing database or transaction log dumps often enough, or you may not have enough space allocated to the database or transaction log. Check for any occurrences of `dump transaction with no_log` in any of your scripts or procedures and replace them with `dump transaction with truncate_only`.

For more information about `dump transaction with no_log`, refer to “dump transaction” in the *Reference Manual: Commands* and “Backing Up and Restoring User Databases” in the Adaptive Server Enterprise configuration guide for your platform. Also refer to “Using the special dump transaction options” in the chapter “Developing a Backup and Recovery Plan” in the *System Administration Guide: Volume 2*.

Warning! Use `dump transaction with no_log` only when there are no active users in the database.

Version in which this error is raised

All versions

Distributed Database Network Errors (7200s)

This section contains error messages for distributed database network code.

Error	Page
Error 7201	523
Error 7205	524
Error 7207	526
Error 7211	527
Error 7212	528
Error 7214	529
Error 7215	530
Error 7218	532
Error 7220	532
Error 7221	535
Error 7223	537
Error 7227	538
Error 7234	539
Error 7235	540

Error 7201

Severity	17
Message text	Can't open a connection to site '%S_SRVID' because there are no remote sites available. Rerun when there are fewer active remote sites, or contact a user with System Administrator (SA) role.
Explanation	The number of remote sites parameter (called remote sites in pre-11.0 releases) determines the number of site connection handlers that can be active at any one time. Site handlers maintain connections between Adaptive Servers, Backup Servers, and Open Server applications.

Error 7201 is raised if there are no site handlers available when the server attempts to open a connection to another server (named in the error text). This occurs when the number of remote entries in `syservers` exceeds the number of remote sites. For example, the error may be seen when you execute a `dump` command, since this requires communication with Backup Server, which in turn requires a site handler.

Action

7201 errors are corrected by configuring sufficient remote sites.

- 1 Check the number of remote sites:

```
1> sp_configure "number of remote sites"  
2> go
```

The `Run Value` in the output shows the number of remote sites that can access the server simultaneously.

- 2 Increase the value of `number of remote sites` as necessary. This is a static parameter, so you must shut down and restart the server to have Adaptive Server use the new value.

Additional information

See the chapter “Managing Remote Servers” in the *System Administration Guide: Volume 1* for more information.

Version in which this error is raised

All versions

Error 7205

Severity

18

Message text

```
Can't open a connection to site '%S_SRVID'. See the  
error log file in the Adaptive Server boot directory.
```

Explanation

This error occurs when Adaptive Server is unable to execute a remote procedure call (RPC) on a remote Adaptive Server, Backup Server, or Open Server application.

Note In this write-up, “Server” is used to indicate an Adaptive Server, Backup Server, or Open Server application.

Action

To determine why this error occurred, perform the following steps:

- 1 Check whether the remote Server is running by issuing an `isql` statement similar to the following:


```
% isql -S <name_of_remote_server>
```

If you receive the following message, the remote Server may not be running.

CT-LIBRARY error:

```
ct_connect(): network packet layer: internal net library error: Net-Lib
protocol driver call to connect two endpoints failed
```

Log in to the remote Server machine and check whether the Server is running. If the Server is not running, start it. Check whether you receive the following message after issuing the `isql` statement:

```
The context allocation routine failed when it tried to load localization
files!!
```

One or more following problems may caused the failure:

```
Your sybase home directory is /export/home/Products/ase1502_64/ASE-15_0.
Check the environment variable SYBASE if it is not the one.
```

```
Cannot access file /export/home/Products/ase1502_64/ASE-
15_0/config/objectid.dat
```

An error occurred when attempting to allocate localization-related structures.

If you do, this indicates that your `SYBASE` environment variable is not set to the correct directory.

The value of the `SYBASE` parameter has to be set to specify the full path name of the `SYBASE` home directory where the *interfaces* (or *sql.ini*) file can be found. The value of the `SYBASE` parameter specifies the device as well as the path for `SYBASE` home. If you move the `SYBASE` directory structure to a new device, make sure the device is included in the `SYBASE` environment variable string.

- 2 Check the error log for both the sending Server and the receiving Server. If other errors exist, refer to the write-ups in this manual for information about recovering from them.
- 3 If, after resolving any problems found in steps 1 and 2, you still get error 7205, call Sybase Technical Support.

Additional information

For more information about managing remote Adaptive Servers, refer to the chapter “Managing Remote Servers” in the *System Administration Guide: Volume 1*.

Before calling Sybase Technical Support, have the information on hand listed in “Reporting errors” on page 4, including the output from `sp_helpserver` and `sp_configure` and the contents of the *interfaces* file.

Version in which this error is raised

All versions

Error 7207

Severity

17

Message text

```
Can't open a connection to site '%S_SRVID' because there are no remote connections available. Rerun when there are fewer remote users and/or fewer local users making remote accesses, or contact a user with System Administrator (SA) role.
```

Explanation

This error occurs when Adaptive Server is unable to open a connection to a remote Adaptive Server, Backup Server, or Open Server application because no remote connections are available.

There may be too many local users with active remote connections. Your System Administrator may also need to configure the server to allow more remote connections.

Action

1 Use the following query to check the remote connections currently in use:

```
1> use master
2> go
```

```
1> select count(*) from sysprocesses a, syssservers b
2> where a.hostname=b.srvname
```

2 Perform one of the following actions to resolve the problem:

- Wait until some of the presently active connections shut down, then retry your query.
- Configure Adaptive Server to allow more remote connections. The relevant configuration parameters are:
 - number of remote connections
 - number of remote logins

These are static parameters. You must restart Adaptive Server after changing the values.

Additional information	Refer to the <i>System Administration Guide: Volume 1</i> for more information. Note that each simultaneous connection to and XP Server for extended stored procedure execution uses up to one remote connection.
Version in which this error is raised	All versions

Error 7211

Severity	18
Message text	Can't open a connection to site '%S_SRVID' because 'set up channel' message was not acknowledged.
Explanation	<p>This error occurs when Adaptive Server is unable to execute a remote procedure call (RPC) on a remote Adaptive Server, Backup Server, or Open Server application because the remote Server did not acknowledge the connection message sent by the local Server.</p> <p>One reason for this error is that the physical connection has died before a response was received. The transaction is then aborted.</p> <p>It may also occur if two access attempts occur within such a short time frame that the first physical connection established has not yet timed out.</p> <hr/> <p>Note In this write-up, “Server” is used to indicate an Adaptive Server, Backup Server, or Open Server application.</p> <hr/>
Action	<p>To find out why this error is occurring, perform the following steps:</p> <ol style="list-style-type: none"> 1 When automatic timeouts are enabled, a timeout occurs when there is no activity for one minute. If you believe timeouts are occurring, consider disabling automatic timeouts by issuing this command: <pre>1> sp_serveroption <server_name> "timeouts", false 2> go</pre> <p>Where <i><server_name></i> is the name of the Adaptive Server for which you are disabling automatic timeouts.</p> 2 Check the error log for both the sending server and the receiving server. If other errors exist, refer to the write-ups in this manual for information about recovering from them.

- 3 If, after resolving any problems found in steps 1 and 2, you still get error 7211, call Sybase Technical Support.

Additional information

For more information about managing remote Adaptive Servers, refer to the chapter “Managing Remote Servers” in the *System Administration Guide: Volume 1*.

Before calling Sybase Technical Support, have the information on hand listed in “Reporting errors” on page 4, including the output for `sp_helpserver` and `sp_configure` and the contents of the *interfaces* file.

Version in which this error is raised

All versions

Error 7212

Severity

11

Message text

Can't close connection <%d> on site <'%S_SRVID'> because it is not active.

Explanation

This error occurs when Adaptive Server is unable to close a connection for a remote site on an Adaptive Server, Backup Server, or Open Server application because the remote connection is not active. <%d> is the logical connection channel number and <%S_SRVID> is the remote Adaptive Server name.

One reason for this error is that the physical connection died before a response was received. The transaction is then aborted.

Error 7212 occurs with these states:

State	Meaning
2	Occurs when Adaptive Server attempts to close the logical connection to the remote site and the logical connection is not found.
3	Occurs when Adaptive Server sends a attention packet to the remote connection and the logical connection is not found.
4	Occurs when Adaptive Server sends a attention packet to the remote connection but the remote site does not exist.

Action

To find out why this error occurred, perform these steps:

- 1 When automatic timeouts are enabled, a timeout occurs when there is no activity from a remote user for one minute. If you believe timeouts are occurring, consider disabling automatic timeouts by issuing the following command:

```
1> sp_serveroption <server_name>, "timeouts", false
2> go
```

Where *<server_name>* is the name of the Adaptive Server for which you are disabling automatic timeouts.

- 2 Check the error log for both the sending Server and the receiving Server. If other errors have been raised, refer to the write-ups in this manual for information about recovering from them.
- 3 If, after resolving any problems found in steps 1 and 2, and error 7212 persists, call Sybase Technical Support.

Additional information

For more information about managing remote Adaptive Servers, refer to the chapter “Managing Remote Servers” in the *System Administration Guide: Volume 1*.

Before calling Sybase Technical Support, have the information on hand listed in “Reporting errors” on page 4, including the output from `sp_helpserver` and `sp_configure` and the contents of the *interfaces* file.

Version in which this error is raised

All versions

Error 7214

Severity

11

Message text

No site name received in server login packet. Server login rejected.

Explanation

This error occurs when executing a remote procedure call (RPC) from an Adaptive Server that has not been given a name using the `sp_addserver` system procedure. The Adaptive Server name referred to here is stored in the `syssservers` system table as well as in the `@@servername` global variable.

Action

- 1 Use `sp_addserver` to assign a name to the Adaptive Server initiating the RPC, while using that server:

```
1> sp_addserver <sql_server_name>, local
2> go
```

Refer to “`sp_addserver`” in the *Reference Manual: Procedures* for information about valid Adaptive Server names.

- 2 Restart Adaptive Server to have the change take effect.

- 3 Verify that the server initiating the RPC has been given a name:

```
1> select @@servername
2> go
```

You can also check the Adaptive Server start message. The output will show the (local) Adaptive Server initiating the RPC. If the output is NULL, then the local server has not yet been assigned a name.

Additional information

For more information, refer to:

- “Server is unnamed” on page 685
- “Identifiers” in the *Transact-SQL User’s Guide* chapter “SQL Building Blocks”

Version in which this error is raised

All versions

Error 7215

Severity

18

Message text

Can't send to site '%S_SRVID'.

Explanation

This error occurs when you try to execute a remote procedure call (RPC) on a remote Adaptive Server or Backup Server that is found in the local interfaces file but is not currently running.

Note In this write-up, “Server” is used to indicate an Adaptive Server or a Backup Server.

Action

- 1 Check whether the remote Server is running by issuing an `isql` statement similar to the following:

```
% isql -S<name_of_remote_server>
```

If you receive this message, then the remote Server may not be running. Log in to the remote Server machine and check whether the Server is running. If it is not running, start it:

```
CT-LIBRARY error: ct_connect(): network packet layer: internal net
library error: Net-Lib protocol driver call to connect two endpoints
failed
```

- 2 Check whether you receive the following message after issuing the `isql` statement:

```
The context allocation routine failed when it tried to load localization
files!!
```

```
One or more following problems may caused the failure
```

```
Your sybase home directory is /export/home/Products/ase1502_64/ASE-15_0.
Check the environment variable SYBASE if it is not the one
```

```
Cannot access file /export/home/Products/ase1502_64/ASE-
15_0/config/objectid.dat
```

```
An error occurred when attempting to allocate localization-related
structures.
```

If you do, this indicates that your *SYBASE* environment variable is not set to the correct directory.

The value of the *SYBASE* parameter must be set to the full path name of the *SYBASE* home directory where the *interfaces* (or *sql.ini*) file can be found. The value of the *SYBASE* parameter specifies the device as well as the path for *SYBASE* home. If you move the *SYBASE* directory structure to a new device, make sure the device is included in the *SYBASE* environment variable string.

Note There are changes to the Adaptive Server Enterprise directory structure in version 12.0. Refer to the installation and configuration guides for your platform for details.

Additional information

For more information about managing remote Adaptive Servers, refer to “Managing Remote Servers” in the in the *System Administration Guide: Volume 1*.

Before calling Technical Support, have the information available listed in “Reporting errors” on page 4, including:

- Adaptive Server (and Backup Server if appropriate) error log for both the local and remote Server
- `sp_helpserver` and `sp_configure` output
- Content of *interfaces* file

Version in which this error is raised

All versions

Error 7218

Severity	11
Message text	Site '%.*s' not found in syssservers.
Explanation	This error occurs when the remote Adaptive Server name in a remote procedure call (RPC) is not defined in the <code>syssservers</code> table on the Adaptive Server initiating the RPC.
Action	<p>Add the remote Adaptive Server name to the <code>syssservers</code> table on the Adaptive Server that initiated the RPC by typing the following:</p> <pre>1> sp_addserver <server_name> 2> go</pre> <p>Refer to “<code>sp_addserver</code>” in the <i>Reference Manual: Procedures</i> for information about valid Adaptive Server names.</p> <p>The easiest way to name an Adaptive Server is by assigning it the name stored in the <code>DSQUERY</code> environment variable or the name that is used in the <code>-S</code> option when the server is accessed via <code>isql</code>. For example:</p> <pre>% isql -SPRODUCTION</pre> <p>In the above example, “<code>PRODUCTION</code>” is the Adaptive Server name. Determine whether the Server has been added to the <code>syssservers</code> table by issuing the following command:</p> <pre>1> sp_helpserver 2> go</pre>
Additional information	For more information, refer to <code>sp_addserver</code> and <code>sp_helpserver</code> in the <i>Reference Manual: Procedures</i> , and “ <code>Identifiers</code> ” in the <i>Reference Manual: Building Blocks</i> .
Version in which this error is raised	All versions

Error 7220

Severity	11
Message text	Site '%s' not found in interfaces file.

Explanation When you execute a remote procedure call (RPC) from one Adaptive Server to another, the local interfaces file of the initiating Adaptive Server has to contain information about the remote Adaptive Server.

Error 7220 occurs in one of the following situations:

- When you try to execute an RPC on an Adaptive Server which is not referenced in the local *interfaces* (or *sql.ini*) file. Its name can, however, be found in the *syssservers* table.
- When you attempt to use the *dump* or *load* commands and the Backup Server is not listed in the *interfaces* (or *sql.ini*) file.

Note Throughout this write-up, your network addressing file is referred to as the “*interfaces* file” although the file name on Windows platforms is *sql.ini*.

Action Eliminate the 7220 error by following the instructions that are appropriate to your situation.

The remote Adaptive Server is not in the local *interfaces* file

Use *dsedit* or *dscp* to check the items in the steps below:

- Determine whether the local *interfaces* file has the query entry for the remote Adaptive Server.
 - Determine whether the Adaptive Server name is spelled correctly in the local *interfaces* file and if the tabs are set correctly at the beginning of each row which contains information about the Adaptive Server ports.
 - If you cannot find any information about the remote Adaptive Server in any *interfaces* file, you have two options. You can either:
 - Use a backup copy of the local or remote *interfaces* file, and add the name of the remote Adaptive Server to the local *interfaces* file manually. Add the query entry for the remote Adaptive Server to the local *interfaces* file
- or:
- Look for lines like the following in your remote Adaptive Server error log (the example is from a Solaris system). If you do not have a backup copy of the local *interfaces* file, check the remote Adaptive Server’s error log for the network name and port numbers at the last restart (that is, at the end of the error log file). Refer to the Open Client/Server installation guide for your platform for more information about Adaptive Server ports.

```
... 00:00000:00008:2007/09/13 15:51:58.34 kernel network name sampa,
interface IPv4, address 158.78.41.25, type tcp, port 8915, filter NONE
....
```

Look for the port number, the network type, and the machine name.

Add an entry to the *interfaces* file like the following:

```
ase1502_64
    master tcp ether sampa 8915
    query tcp ether sampa 8915
```

Warning! Before you add Adaptive Server's query, master, and console entries to the local *interfaces* file, verify that no other Adaptive Server has the same port numbers assigned in the local *interfaces* file.

- After changing the *interfaces* file, try accessing the remote Adaptive Server through `isql` or, if it is not running, restart it.
- Also check your hardware error log for possible hardware failures, in case you still cannot reboot or access the remote Adaptive Server.

The Backup Server is not listed in the *interfaces* file

The entry for the local Backup Server is created when you install Adaptive Server. If you have installed a remote Backup Server on another machine, choose one of these options to resolve the error:

- Copy the *interfaces* file entry for this Backup Server to the *interfaces* file on the machine where your Adaptive Server runs.
- Locate the *interfaces* file on a file system that is shared by both machines.
- Use the `dsedit` or `dscp` utility to update your *interfaces* file on the machine where your Adaptive Server runs. Refer to the Adaptive Server configuration and installation guide for your platform for information about these utilities.

Additional information

For more additional information, refer to the “allow remote access” configuration parameter in the *System Administration Guide: Volume 1* chapter “Setting Configuration Parameters,” and “*interfaces* file” in the Adaptive Server Enterprise installation and configuration guide for your platform.

Version in which this error is raised

All versions

Error 7221

Severity	14
Message text	Login to site '%S_SRVID' failed.
Explanation	<p>This error occurs when you execute a remote procedure call (RPC) in one of the following situations:</p> <ul style="list-style-type: none"> • A name was not assigned, via <code>sp_addserver</code>, to the Adaptive Server which is initiating the RPC. • A remote login for the Adaptive Server initiating the RPC was not added to the Adaptive Server that is receiving the RPC. • The Adaptive Server receiving the RPC is not configured for allow remote access or for high enough values for number of remote sites or number of remote connections. • The Adaptive Server receiving the RPC does not have an entry in the network addressing (<i>interfaces</i> or <i>sql.ini</i>) file for the Server initiating the RPC (or the other way around). • The Adaptive Server receiving the RPC is in single-user mode and already had one user at the time the RPC was executed. • The password for the user initiating the RPC is not the same on both Adaptive Servers.
Action	<p>Eliminate the 7221 error by determining which of the above situations applies to you and following the appropriate steps:</p> <ol style="list-style-type: none"> 1 A name was not assigned to the Adaptive Server that is initiating the RPC. <ul style="list-style-type: none"> Use the <code>select @@servername</code> command to check if the Adaptive Server initiating the RPC has a name assigned to it. If the <code>select @@servername</code> command returns the name of the server which is initiating the RPC, then continue to step 2. If the output returned by <code>select @@servername</code> is NULL, then assign a name to the local Adaptive Server using the <code>sp_addserver</code> command: <pre>1> sp_addserver <name_of_server_initiating_rpc>,local 2> go</pre> <p>Refer to “<code>sp_addserver</code>” in the <i>Reference Manual: Procedures</i> for information about valid Adaptive Server names.</p>

Restart Adaptive Server after assigning it a name. Usually you want to name an Adaptive Server the same name as is stored in the *DSQUERY* environment variable or the name which is used in the -S option when the server is accessed via isql. For example:

```
% isql -SPRODUCTION
```

In the above example, “PRODUCTION” is the Adaptive Server name.

- 2 A remote login was not added to the Adaptive Server that is receiving the RPC for the Adaptive Server that is initiating it.

Use `sp_helpremotelogin` to check if the receiving Adaptive Server has a remote login for the Adaptive Server initiating the RPC. If the output returned by `sp_helpremotelogin` contains information about the remote Adaptive Server, then continue to step 3.

If the output returned by `sp_helpremotelogin` does not contain information about the remote Adaptive Server, then add a remote login to the remote Adaptive Server for the Adaptive Server that is initiating the RPC by typing the following (on the remote Adaptive Server):

```
1> sp_addremotelogin <name_of_server_initiating_rpc>
2> go
```

- 3 The Adaptive Server receiving the RPC is not configured for allow remote access or for high enough values for number of remote sites or number of remote connections.

Check the remote Adaptive Server’s error log for more information on why the RPC initiated by the local Adaptive Server did not succeed. Refer to “Setting Configuration Parameters” in the *System Administration Guide: Volume 1* for information about how to turn on allow remote access and how use `sp_configure` to raise the values of number of remote sites or number of remote connections.

- 4 The Adaptive Server receiving the RPC does not have an entry in the *interfaces* file for the Server initiating the RPC (or the other way around).

Add the required entries for the Adaptive Server initiating the RPC to the *interfaces* file of the receiving Server. Use one of the following utilities or copy the entries from the *interfaces* file of the initiating Server and make sure that they have a tab in front of each entry.:

- dsedit
- dscp

5 The Adaptive Server receiving the RPC is in single-user mode and already has one user at the time the RPC is executed.

Wait until the receiving Adaptive Server is not being used, then execute the RPC or restart the receiving Adaptive Server in multiuser mode.

6 The password for the user initiating the RPC is not the same on both Servers.

Make user passwords the same for the remote Adaptive Server.

Additional information

For more information, refer to:

- `sp_addremotelogin`, `sp_addserver`, and `sp_configure` in the *Reference Manual: Procedures*
- “Server is unnamed” on page 685 of this manual
- “Allowing remote access” in *System Administration Guide: Volume 1*, in the chapter “Managing Remote Servers.”

Version in which this error is raised

All versions

Error 7223

Severity

16

Message text

Cannot %S_MSG site '%S_SRVID'.

Explanation

This error occurs when Adaptive Server fails to complete a read or send operation during a remote procedure call (RPC). The failure can occur in a number of states. Each of these states is described below.

Note Adaptive Server uses tokens to establish a dialog request with a client. There are several kinds of tokens to send different signals.

State	Meaning
1	Adaptive Server uses buffering for sending and receiving packets to and from remote connections. If an error occurs while Adaptive Server is trying to read the next packet from the remote connection while Adaptive Server is being run in debug mode, error 7223 is raised with state 1. %S_MSG is replaced by “read from” and %S_SRVID is the remote Adaptive Server name.

State	Meaning
2	If Adaptive Server confirms that no token was read from the remote connection, then error 7223 is raised with state 2. <i>%S_MSG</i> is replaced by “read from” and <i>%S_SRVID</i> is the remote Adaptive Server name.
3	If an error occurs while Adaptive Server is sending a token to a remote connection, error 7223 is raised with state 3. <i>%S_MSG</i> is replaced by “send to” and <i>%S_SRVID</i> is the remote Adaptive Server name.
4	If an error occurs while Adaptive Server is sending data to a remote connection, error 7223 is raised with state 4. <i>%S_MSG</i> is replaced by “send to” and <i>%S_SRVID</i> is the remote Adaptive Server name.
5	While Adaptive Server is establishing a network connection to a remote Adaptive Server, it translates the remote Server name to a Server ID, opens the logical connection, and then logs in. Error 7223 is raised with state 5 when Adaptive Server fails during one of these operations. <i>%S_MSG</i> is replaced by “send to” and <i>%S_SRVID</i> is the remote Adaptive Server name.
6	Adaptive Server builds a remote procedure call request message in the remote connection send buffer and then sends it. Error 7223 occurs with state 6 when this send fails. <i>%S_MSG</i> is replaced by “send to” and <i>%S_SRVID</i> is the remote Adaptive Server name.
7	If Adaptive Server confirms that no data was read from the remote connection, then error 7223 is raised with state 7. <i>%S_MSG</i> is replaced by “read from” and <i>%S_SRVID</i> is the remote Adaptive Server name.
12	A connection to a remote server using CIS has failed.
13	While establishing a remote connection, a “send to” the remote server failed.
14	While responding to a security request from a remote server, the “send” failed.
15	While responding to an encrypted password request, the “send” failed.

Action	<ol style="list-style-type: none"> 1 Check the network connection between the locations of your local and remote Adaptive Servers. Make sure packets of data can be sent and received. You may need to ask your network administrator for help. 2 If you did not find a problem in step 1, call Sybase Technical Support.
Additional information	<p>When you call Sybase Technical Support, have the information ready that is listed in “Reporting errors” on page 4, including:</p> <ul style="list-style-type: none"> • Server error log for both the local and remote Adaptive Servers • <code>sp_configure output</code>
Version in which this error is raised	All versions

Error 7227

Severity 17

Message text	Can't open a connection to site '%S_SRVID' because Adaptive Server is not configured for remote access. Ask a user with System Administrator (SA) role to reconfigure Adaptive Server to allow remote access.
Explanation	This error occurs when you execute a remote procedure call (RPC) and the sending or receiving server is not configured for remote access.
Action	<p>Configure the local or remote Adaptive Server for remote access by changing the allow remote access configuration parameter via <code>sp_configure</code>.</p> <pre>1> sp_configure "allow remote access", 1 2> go</pre> <p>After enabling the allow remote access configuration parameter and restarting Adaptive Server, the number of remote logins, number of remote sites, number of remote connections, and remote server pre-read packets configuration parameters are set to their default values.</p> <p>If you determine that the default values for number of remote logins, number of remote sites, number of remote connections, and remote server pre-read packets are not adequate for your needs, use <code>sp_configure</code> to change these configuration parameters as well.</p> <p>Restart Adaptive Server to make the configuration change take effect.</p>
Additional information	Refer to “Allowing remote access” in <i>System Administration Guide: Volume 1</i> and “ <code>sp_configure</code> ” in the <i>Reference Manual: Procedures</i> .
Version in which this error is raised	All versions

Error 7234

Severity	10
Message text	Failed to connect to remote server %s %s %s %s %s for site %s. Trying the next query entry of this server, if any.
Explanation	Communication between an Adaptive Server and another server requires a site handler to manage connections between the servers. When you log on to Server A, and issue a command that necessitates communication with Server B (such as Backup Server for dumps and loads), Server A uses a site handler to open a connection to Server B, sending the query and receiving results from Server B.

Error 7234 is raised when Adaptive Server is unable to open a connection to another server. The `remote server` information in the message text refers to the:

- Server name
- Network type
- Network name
- Machine name
- Port number

Processing continues after the error is raised, and the server retries the connection using the next entry, if any, in the *interfaces* file.

Error 7234 is an informational message. However, if all attempts to open a connection using the available *interfaces* entries fail, Adaptive Server raises Error 7235 ("No other query entry found for the remote server %s - site %s. Aborting this site handler, as no remote server responded. Please make sure the remote server is accepting connections.").

Action

- Verify that the remote server is up and running. If you see the error after issuing a dump command, for example, check that Backup Server is running.
- Verify the remote server's network name in `syssservers`:

```
1> select srvname, srvnetname from master..syssservers
2> go
```

The remote server's `srvnetname` entry must match the name appearing in the *interfaces* file.

Version in which this error is raised

All versions

Error 7235

Severity

17

Message text

```
No other query entry found for the remote server %s -
site %s. Aborting this site handler, as no remote server
responded. Please make sure the remote server is
accepting connections.
```


Explanation	<p>Communication between an Adaptive Server and another server requires a site handler to manage connections between the servers. When you log on to Server A, and issue a command that necessitates communication with Server B (such as Backup Server for dumps and loads), Server A uses a site handler to open a connection to Server B, sending the query and receiving results from Server B.</p> <p>When the first attempt to open a site handler connection fails, Adaptive Server raises Error 7234 ("Failed to connect to remote server %s %s %s %s %s for site %s. Trying the next query entry of this server, if any."). The server then re-tries the connection using the next entry, if any, in the <i>interfaces</i> file. Error 7235 is raised after all such connection attempts fail.</p>
Action	<ul style="list-style-type: none"> • Check that the remote server is up and running. If you see the error after issuing a dump command, for example, type the following: <pre style="margin-left: 40px;">1> exec SYB_BACKUP...sp_who 2> go</pre> <p>If Backup Server is running, you will get a meaningful response to this query.</p> • Verify the remote server's network name in <i>syssservers</i>: <pre style="margin-left: 40px;">1> select srvname, srvnetname from master..syssservers 2> go</pre> <p>The remote server's <i>srvnetname</i> entry must match the name appearing in the <i>interfaces</i> file. (An incorrect <i>srvnetname</i> can prevent the server from reaching the correct remote server.)</p>
Additional information	When you call Sybase Technical Support, have the information ready that is listed in "Reporting errors" on page 4.
Version in which this error is raised	All versions

Text Manager Errors (7100s)

This section contains error messages for the Adaptive Server Enterprise Text Manager.

Error	Page
Error 7101	543
Error 7105	544
Error 7114	546
Error 7130	547
Error 7134	549

Error 7101

Severity

17

Message text

Unable to allocate new text value, dbid %d.

Explanation

When you create a table with text, image, or unitext columns, the data is stored on its own chain of text, image or unitext pages. Error 7101 is raised when you attempt to insert data into a text, image, or unitext column in a table, but there is not enough space to allocate the new text, image, or unitext pages.

Error 7101 is raised with the following states:

State	Meaning
1	Could not allocate the first page to add a new value in a table's text, image, or unitext column.
2	Could not allocate additional pages for an in-progress insert to a text, image, or unitext column.

Error 7101 is similar to error 1105, states 1 and 2, but it is specific to text, image, or unitext allocation.

Action

1 Check the available space in the database:

```
1> use <database_name>
```

```
2> go
1> sp_spaceused
2> go
```

- 2 Identify the table by examining the query that raised the error. (You can also run `sp_helpsegment` on the database segments; rows where `indid = 255` refer to the tables with text, image, or unitext columns.)

- 3 Check the space use in the table:

```
1> sp_spaceused <table_name>, 1
2> go
```

To find the space used by text, image, or unitext columns in this report, look for objects named “t” plus the table name.

- 4 Use `alter database` to increase the storage space for the text, image, or unitext data.

Additional information

By default, a chain of text, image, or unitext pages is placed on the same device as its table. You can specify a separate device for future space allocations to the text, image, or unitext pages by using the `sp_placeobject` system stored procedure. This is desirable when you expect substantial growth in text, image, or unitext storage.

See “Placing Text Pages on a Separate Device” in the chapter “Creating and Using Segments” in the *System Administration Guide: Volume 2* for details.

Version in which this error is raised

All versions

Error 7105

Severity

22, 11

Message text

```
TEXT/IMAGE/UNITEXT page %d does not have a next page,
although it should.
```

Note This error may be caused by a hardware problem.

Explanation

Tables can contain text, image, or unitext columns. Adaptive Server stores text, image, or unitext data in a linked list of data pages separate from the rest of the table. On a 2K page size server, each 2K text, image, or unitext page stores a maximum of 1800 bytes of data.

Error 7105 occurs when Adaptive Server attempts data retrieval or another operation on a table's text, image, or unitext data, and the operation fails due to an unexpected termination of the page chain. If the operation involved copying data, the problem could be on the source side (the page chain being read) or the destination side (the page chain being written out). The severity of the 7105 error message can help you determine the origin of the problem.

Error 7105 indicates a bad text, image, or unitext page chain; when the severity of the error is 11, the error may result from doing dirty reads.

Error 2574 is a related error (see "Error 2574" on page 325); it is raised when a dbcc check finds problems with a table's text, image, or unitext page chain.

Action

Severity 22:

The problem is typically due to a small number of bad rows in a table with text, image, or unitext data, and can be corrected by dropping and re-creating those rows. Use the following steps to correct the problem:

- 1 Identify the table. Refer to "Finding an object name from a page number" in the chapter "Other Useful Tasks" in the most recent version of *Troubleshooting and Disaster Recovery*.

Warning! The procedure referenced in step 1 uses the dbcc page command. Use the dbcc page command only as directed. This command is undocumented, nonstandard, and is provided "as is" without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

- 2 Determine the first page of the text, image, or unitext chain using dbcc pgl linkage, with the starting page being the page number from the error message.

For example, assuming dbid=7 and starting page number 531:

```
1> dbcc pgl linkage(7, 531, 0, 1, 0, 0)
2> go
Object ID for pages in this chain = 208003772.
Beginning of chain reached.
Page : 531
Page : 532
Page : 530
3 pages scanned. Object ID = 208003772.
Last page in scan = 530.
```

Page 530 is the first page of the text, image, or unitext chain.

- 3 Identify the rows with the bad data using the following command:

```
1> select <column1> from <table_name>
2> where convert (int, textptr(<text_colname>))
3> = <first_page_of_chain>
4> go
```

Where *<column1>* is any column or columns in your table that will uniquely identify the rows.

- 4 Drop the rows in question.
- 5 Insert the rows back into the table.
- 6 Run `dbcc checktable` on the table to verify that the problem has been corrected.

Severity 11:

Error 7105 is probably due to doing a dirty read, and is unlikely to be related to page corruption. If the problem persists on the same page, page corruption may be involved; follow the instructions for severity 22 (above) to resolve the problem.

Additional information

When you call Sybase Technical Support, have the information on hand that is listed in “Reporting errors” on page 4, including the output of `dbcc page` and `dbcc pglinkage`.

Version in which this error is raised

All versions

Error 7114

Severity

22

Message text

Page %ld is not a valid text page.

Note This error may be caused by a hardware problem.

Explanation

Error 7114 occurs when Adaptive Server attempts to access or update a table’s text or image data, but the retrieved page is found to be a non-text page.

Text page corruption can give rise to 605 errors ("An attempt was made to fetch logical page '%ld' in database '%.*s' from cache '%.*s'. Page belongs to object with id '%ld', not to object '%.*s'.) when you run dbcc's on the affected data.

This error may be due to hardware or operating system errors, or an Adaptive Server problem.

Action	<p>Contact Sybase Technical Support for assistance when this error is raised. Before calling, collect relevant information using the following steps:</p> <ol style="list-style-type: none"> 1 Note the page number in the error message. Using the procedure in “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of <i>Troubleshooting and Disaster Recovery</i> to identify the type of page and the table associated with the page. 2 Check the integrity of the table’s text page chains: <pre> 1> dbcc textalloc(<tablename>) 2> go </pre>
Additional information	When you call Sybase Technical Support, have the information on hand that is listed in “Reporting errors” on page 4, including the output of dbcc checks:
Version in which this error is raised	All versions

Error 7130

Severity	16
Message text	<pre> WRITETEXT with no log is not valid at this time. Database %s must first be configured for this. </pre>
Explanation	<p>This error occurs when you try to execute the writetext command or dbwritetext routine in a database which does not have the sp_dboption option select into/bulkcopy enabled. Because these operations do not log the changes to your data, the option must be enabled.</p> <p>Often, writetext and dbwritetext run against a temporary table. The select into/bulkcopy option must be enabled in tempdb before these will run successfully.</p>

Action	<p>Enable the <code>sp_dboption</code> option <code>select into/bulkcopy</code> on all databases affected by the <code>writetext</code> just before you perform this nonlogged operation. To do this, the System Administrator or Database Owner can use the following steps:</p> <ol style="list-style-type: none">1 Enable the option using this procedure:<pre>1> use master 2> go 1> sp_dboption <database_name>, 2> "select into/bulkcopy", true 3> go</pre>2 To verify that the change is active, enter:<pre>1> sp_helpdb <database_name> 2> go</pre>3 Execute the nonlogged operation.4 When the nonlogged operation is complete, dump the database using normal system administration procedures.5 Return the database to its original condition by disabling the <code>select into/bulkcopy</code> option:<pre>1> use master 2> go 1> sp_dboption <database_name>, 2> "select into/bulkcopy", false 3> go</pre> <p>Repeat step 2 above if you want to verify that the change is active.</p>
Additional information	<p>Refer to “<code>writetext</code>” in the <i>Reference Manual: Commands</i> and to “<code>dbwritetext</code>” in the “Routines” chapter of the <i>Open Client DB-Library Reference Manual</i>.</p> <p>For more information about database dumps and recoverability, refer to “Developing a Backup and Recovery Plan” in the <i>System Administration Guide: Volume 2</i>.</p>
Version in which this error is raised	All versions

Error 7134

Severity	16
Message text	The text table and the table referenced by the text pointer value %s disagree. This text pointer references first text page %d, which belongs to object ID %d.
Explanation	<p>Adaptive Server stores text, image, or unitext data in text extents separate from the rest of the table. Each extent contains text data for a given table, and may contain multiple text chains corresponding to text contained in different rows and text columns of the table.</p> <p>The Adaptive Server Text Manager validates text, image, or unitext pages during processing. Error 7134 is raised if the Text Manager finds that the object ID of a text, image, or unitext page does not match the object ID of the text, image, or unitext extent to which the page belongs.</p>
Action	<p>If error 7134 is raised when you select text, image, or unitext columns within a trigger, note that Adaptive Server does not provide the ability to select text, image, or unitext datatypes from the inserted or deleted table. You must re-code the trigger to resolve this problem.</p> <p>If the error is raised in some other context not involving triggers, the problem is typically due to a small number of bad rows in a table with text, image, or unitext data. To correct the problem, use the following procedure to drop and re-create those rows.</p> <ol style="list-style-type: none"> 1 Identify the table for which the error was raised. 2 Run <code>dbcc checktable</code> on the table. 3 Correct any errors detected by <code>dbcc checktable</code> using the sections in this manual that refer to those errors. 4 If the error persists, contact Sybase Technical Support for assistance.
Additional information	When calling Technical Support, have the information available that is listed in “Reporting errors” on page 4.
Version in which this error is raised	All versions

Threshold Manager Errors (7400s)

This chapters error messages pertain to Adaptive Server Enterprise thresholds. Refer to the chapter “Managing Free Space with Thresholds” in the *System Administration Guide: Volume 2* for information about thresholds.

Error	Page
Error 7401	551
Error 7402	553
Error 7403	554
Error 7404	555
Error 7406	556
Error 7407	557
Error 7408	558
Error 7409	559
Error 7410	560
Error 7411	561
Error 7412	562
Error 7413	563
Error 7414	563
Error 7415	564
Error 7417	565
Error 7418	566
Error 7438	567

Error 7401

Severity

17

Message text

Could not allocate memory for the threshold task's argument list; threshold task cannot be spawned for database %d, segment %d, free space %ld.

Explanation	This error occurs when a stored procedure associated with a threshold was not initiated because there was not enough memory available in procedure cache.
Action	Although error 7401 indicates that the threshold task did not run because not enough memory was available in procedure cache at that time, the threshold task may have run later when enough memory did become available.

Note Before performing the following steps, determine whether the threshold task ran at a later time. If the threshold task has already run, the following steps are not necessary.

To perform the threshold task that did not run because of this error:

- 1 Select the database:

```
1> use <database_name>
2> go
```

Where *<database_name>* is the database for which the error occurred.

- 2 Run `sp_helpthreshold <segment_name>` for the segment named in the error message. This will display the stored procedure associated with the threshold.
- 3 Run the stored procedure manually.

To prevent this error from occurring in the future, make sure there is enough memory to run stored procedures associated with a threshold by testing them during the time when your system is most heavily loaded. If there is not enough memory, do one of the following:

- Use `sp_configure` to increase the Adaptive Server procedure cache size configuration parameter (increases the percentage of total cache space that is used as procedure cache).
- Use `sp_configure` to increase the value of the total memory configuration parameter.
- Rewrite the stored procedure to execute in the procedure cache available for the Adaptive Server.

Additional information	Refer to the Adaptive Server Enterprise <i>Performance and Tuning</i> guides and “Configuring Memory” in the <i>System Administration Guide: Volume 2</i> for information about Adaptive Server memory use.
------------------------	---

Refer to the chapter “Managing Free Space with Thresholds” in the *System Administration Guide: Volume 2* for information about Adaptive Server thresholds.

Version in which this error is raised All versions

Error 7402

Severity	10
Message text	Could not spawn the threshold handler procedure for database %d, segment %d, when only %ld free pages were left in the segment.
Explanation	Thresholds access Adaptive Server as a detached user. This means that, although a threshold does not access Adaptive Server through the network, it still requires a user connection. Error 7402 occurs when there are no available user connections when an Adaptive Server process reaches a threshold and Adaptive Server tries to execute the associated threshold procedure.
Action	<p>To perform the threshold task that did not run because of this error:</p> <ol style="list-style-type: none"> 1 Select the database: <ul style="list-style-type: none"> 1> use <database_name> 2> go <p>Where <database_name> is the database for which the error occurred.</p> 2 Run <code>sp_helpthreshold <segment_name></code> for the segment named in the error message. This displays the stored procedure associated with the threshold. 3 Run the stored procedure manually. <p>To prevent this error from occurring in the future, use <code>sp_configure</code> to make sure that there are enough user connections to allow for all clients and the threshold manager. Refer to the chapter “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for information about changing the number of user connections configuration parameter.</p>
Additional information	Refer to “Managing Free Space with Thresholds” in the <i>System Administration Guide: Volume 2</i> for information about Adaptive Server thresholds.
Version in which this error is raised	All versions

Error 7403

Severity	10
Message text	Threshold task could not use database %d, and so cannot execute the threshold procedure for segment %d, free space %ld.
Explanation	<p>A threshold is “owned” by the user who creates it, or the last user to modify it with <code>sp_modifythreshold</code> (refer to <code>sp_addthreshold</code> and <code>sp_modifythreshold</code> in the <i>Reference Manual: Procedures</i> for details). Error 7403 occurs when the Threshold Manager tries to use a database and execute a stored procedure and access is denied. The procedure is not executed.</p>
Action	<p>To perform the threshold task that did not run because of this error:</p> <ol style="list-style-type: none">1 Select the database:<pre>1> use <database_name> 2> go</pre><p>Where <code><database_name></code> is the database for which the error occurred.</p>2 Run <code>sp_helpthreshold <segment_name></code> for the segment named in the error message. This will display the stored procedure associated with the threshold.3 Run the stored procedure manually using the account of the owning user. To determine the owner of the procedure and user access to the database, type:<pre>1> sp_help <procedure_name> 2> go 1> sp_helpuser <database_name> 2> go 1> sp_helprotect <procedure_name> 2> go</pre> <p>To prevent this error from occurring in the future, make sure that the owner of each threshold has access to the database to which the threshold applies, and that the user has execute permission on the stored procedure associated with the threshold.</p>
Version in which this error is raised	All versions

Error 7404

Severity	10
Message text	Could not find row in <code>systhresholds</code> for database <code>%. *s</code> describing segment <code>%d</code> , free space <code>%ld</code> .
Explanation	<p>This error can occur for either of the following reasons:</p> <ul style="list-style-type: none"> • The row in <code>database.dbo.systhresholds</code> describing the segment, free space, and procedure name for the segment named in the error message is incorrect. Therefore, the threshold manager was not able to execute the stored procedure associated with the threshold. • The clustered index on <code>systhresholds</code> is corrupt and Adaptive Server was not able to retrieve the row.
Action	<p>Run <code>dbcc checktable (systhresholds)</code> and <code>dbcc tablealloc (systhresholds)</code> to check consistency of the system table. If you have corruption on <code>systhresholds</code>, contact Sybase Technical Support for further instructions.</p> <p>Perform the threshold task that did not run because of this error:</p> <ol style="list-style-type: none"> 1 Select the database: <pre>1> use <database_name> 2> go</pre> <p>Where <code><database_name></code> is the database for which the error occurred.</p> 2 Run <code>sp_helpthreshold <segment_name></code> for the segment named in the error message. This will display the stored procedure associated with the threshold. 3 Run the stored procedure manually. <p>If the stored procedure is not found or there are incorrect entries for segment or free space, modify the row in <code>systhresholds</code> with <code>sp_modifythreshold</code> to add a stored procedure (threshold action). Refer to the <i>Reference Manual: Procedures</i> for information about <code>sp_modifythreshold</code>.</p>
Additional information	<p>When calling Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the output of:</p> <ul style="list-style-type: none"> • <code>select @@version</code> • All error messages • <code>dbcc</code> commands • <code>sp_helpthreshold</code>

Version in which this error is raised All versions

Error 7406

Severity	16
Message text	%s: unknown option: %.*s
Explanation	lct_admin is a function used to manage the last-chance threshold (LCT) on the log segment of a database. Error 7406 occurs when an invalid option is passed to lct_admin.
Action	<p>Use one of the valid options listed below, which are all strings and require quotation marks.</p> <ul style="list-style-type: none"> • lct_admin ("reserve", <number_of_log_pages>) – returns the number of pages required to do a successful dump transaction given the number of pages in a log segment. If the provided number of log pages is 0 (zero), the lct_admin() function returns the last-chance threshold value set for the current database. • lct_admin ("lastchance", <database_ID>) – creates a last-chance threshold for a database with log on a separate device. This can be used to create a last-chance threshold for databases that were upgraded from a pre-version 10.0 server. <hr/> <p>Note You should not have to use the lastchance option of lct_admin.</p> <hr/> <ul style="list-style-type: none"> • lct_admin ("logfull", <database_ID>) – determines if the last-chance threshold has been reached in a specified database. The procedure will return “0” if last-chance has not been reached, and “1” if it has.

- `lct_admin("unsuspend", <database_ID>)` – this command is no longer supported. To abort the tasks sleeping on a log-suspend state, consider using `lct_admin(abort)`.

Note In addition to the `lct_admin` options listed above, the following `lct_admin` options also exist, but are not fully documented in the current version of this guide. Contact Sybase Technical Support for more information.

- `lct_admin(abort)`
 - `lct_admin(chkfreespace)`
 - `lct_admin(logsegment_freepages)`
 - `lct_admin(num_logpages)`
 - `lct_admin(reserved_for_rollbacks)`
 - `lct_admin(status_in_cache)`
-

Additional information

Refer to the *Reference Manual: Building Blocks* in the section “String functions” for a description of valid `lct_admin` options for your version.

Version in which this error is raised

All versions

Error 7407

Severity	10
Message text	Failed to execute the threshold procedure for database %d, segment %d, free space %ld.
Explanation	This error occurs when an attempt to execute a stored procedure defined in <code>systhresholds</code> fails.
Action	To perform the threshold task that did not run because of this error: <ol style="list-style-type: none"> 1 Select the database: <pre>1> use <database_name> 2> go</pre> <p>Where <code><database_name></code> is the database for which the error occurred.</p>

- 2 Run `sp_helpthreshold <segment_name>` for the segment named in the error message. This displays the stored procedure associated with the threshold.
- 3 Run the stored procedure manually.

To determine whether you can prevent this error from occurring in the future, review the Adaptive Server error log for additional messages regarding this error. If there is specific information as to what caused the procedure to fail, take whatever steps are necessary to prevent future occurrences.

Version in which this error is raised

All versions

Error 7408

Severity 20

Message text `Could not find a dtable for database %d.`

Explanation Information about each database is kept in memory in a structure called a `dtable`. Error 7408 occurs when an attempt to execute the built-in function `lct_admin` fails because the database table (database ID) information was either not available or incorrect.

Action Verify that you have entered the correct database name (or database ID) in the threshold procedure. If it is not correct, modify the procedure.

This error can occur because the database is not open because the value of the number of open databases configuration parameter is too low. If the value for number of open databases is too low, use `sp_configure` to raise the value.

If the database name is correct, the value of the number of open databases configuration parameter is large enough to accommodate all the system and user databases, and the command still fails with this message, check the Adaptive Server error log for messages.

To perform the threshold task that did not run because of this error:

- 1 Select the database:

```
1> use <database_name>
2> go
```

Where `<database_name>` is the database for which the error occurred.

- 2 Run `sp_helpthreshold <segment_name>` for the segment for which the error occurred. This will display the stored procedure associated with the threshold.
- 3 Run the stored procedure manually.

Version in which this error is raised

All versions

Error 7409

Severity

20

Message text

This task tried to sleep waiting for space to become available in the log, but could not. The current command is being aborted because no further database changes can be logged.

Explanation

The last-chance threshold has been reached in a database. All transactions trying to write to the transaction log should be put to sleep until space is free in the transaction log. Adaptive Server was unable to put the transaction to sleep, so the transaction has been aborted.

Action

Verify that the threshold action defined for the last-chance threshold has executed successfully. If it did not complete, manually execute a `dump transaction` command to create free space in the log. Then rerun the transaction that was aborted. If the transaction still fails, check the Adaptive Server error log for additional information.

Refer to the chapter “Managing Free Space with Thresholds” in the *System Administration Guide: Volume 2* for information about last-chance thresholds.

Version in which this error is raised

All versions

Error 7410

Severity	26
Message text	%s: function called with invalid argument.

Note This error is only raised when using diagserver.

Explanation This error occurs when an attempt to execute a threshold action fails because of a problem in a memory structure.

Action If there are additional messages in the server error log that tell you which threshold was initiated, do the following steps:

- 1 Select the database:

```
1> use <database_name>
2> go
```

Where *<database_name>* is the database which contains the failed threshold.

- 2 Determine the stored procedure that is associated with the segment for which the error occurred:

```
1> sp_helpthreshold <segment_name>
2> go
```

- 3 Run the stored procedure manually.

If you were not able to determine which threshold was initiated by looking at the error log output, run the following `curunreservedpgs` query (appropriate to the version of Adaptive Server Enterprise that you are using) in each database to determine which threshold action did not execute properly:

- Versions prior to 15.0:

```
1> use master
2> go

1> select db_name(dbid), d.name,
2> curunreservedpgs (dbid, lstart, unreservedpgs)
3> from sysusages u, sysdevices d
4> where d.low <= u.size + vstart
5> and d.high >= u.size + vstart - 1
6> and d.status &2 = 2
7> go
```

- Version 15.0 and later

```

1> use master
2> go

1> select db_name(dbid), d.name,
2> curunreservedpgs (dbid, lstart, unreservedpgs)
3> from sysusages u, sysdevices d
4> where d.vdevno = u.vdevno
5> and d.status &2 = 2

```

This query returns the database name, device name, and the number of unreserved pages for each device fragment. If you can determine which threshold was reached (compare the output of the `curunreservedpgs` query with the free space for the segment), execute the threshold action manually. Refer to the chapter “Transact-SQL Functions” in the *Reference Manual: Building Blocks* for information about `curunreservedpgs`.

To prevent this error from occurring in the future, review the Adaptive Server error log and look for other messages associated with this error. There may be other errors that will tell you why the threshold failed. If more information is available, take the necessary steps to prevent the error from recurring.

Version in which this error is raised

All versions

Error 7411

Severity	18
Message text	Table 'systhresholds' in database '%.*s' is not in its correct sort order. Either the clustered index is missing or there is data corruption in the table.
Explanation	<p>Every database in the server has a last-chance threshold. The threshold is an estimate of the number of free log pages required to back up the transaction log. A database's last-chance threshold is stored in the <code>systhresholds</code> table; the LCT and related information is also kept in threshold memory cache.</p> <p>When the server reads the <code>systhresholds</code> data for the database and finds that the information is not in its correct order, it raises Error 7411. There is possible corruption on <code>systhresholds</code> or its clustered index. You may also see error 7417.</p> <p>Error 7411 is raised with the following states:</p>

State	Meaning
1	This state is raised when the server reads the <code>systhreshold</code> table to update the threshold cache in memory for the database, and finds that this information is not in correct order.
2	This state is raised when you run <code>dbcc checkcatalog</code> , the server reads the <code>systhreshold</code> table and finds that the free-space values for the log segments are not in correct order.

Error 7411 indicates a `systhreshold` problem. When accompanied by 7417 and (or) 7418 errors, it may be due to an Adaptive Server problem.

Action

If there are no other errors, run `dbcc checktable (systhresholds)` and `dbcc tablealloc (systhresholds)` to check consistency of the system table. If you find corruption in `systhresholds`, call Sybase Technical Support for further instructions.

If the error is accompanied by 7417 and 7418 errors, you may need to upgrade to the latest Adaptive Server Enterprise version available for your platform.

Additional information

See the write-ups for errors 7417 (“Error 7417” on page 565) and error 7418 (“Error 7418” on page 566).

When you call Sybase Technical Support, have the information listed in “Reporting errors” on page 4 available, including the output of:

- `select @@version`
- `dbcc` commands

Version in which this error is raised

All versions earlier than version 15.0

Error 7412

Severity

10

Message text

Space available in the log segment has fallen critically low in database '%.*s'. All future modifications to this database will be %S_MSG until the log is successfully dumped and space becomes available.

Explanation

This message notifies users that the last-chance threshold has been reached. All transactions requiring access to the transaction log will either be suspended or aborted until the appropriate action is taken to clear the transaction log.

Action	<ol style="list-style-type: none"> 1 Ensure that the action defined by the last-chance threshold procedure was able to execute successfully. If the procedure failed for any reason, manually run <code>dump transaction</code> to create free space in the transaction log. 2 Look for errors in the Adaptive Server error log to determine why the procedure failed. 3 Restart any transactions that have been aborted.
Additional information	Refer to the chapter “Managing Free Space with Thresholds” in the <i>System Administration Guide: Volume 2</i> for information about last-chance thresholds.
Version in which this error is raised	All versions

Error 7413

Severity	10
Message text	<code>%d task(s) are sleeping waiting for space to become available in the log segment for database %.*s.</code>
Explanation	This message is written to the Adaptive Server error log by the checkpoint process when the last-chance threshold has been exceeded, and the specified number of processes have been put to sleep.
Action	Ensure that the appropriate last-chance threshold action is executed to create free space in the log. If the procedure did not run for any reason, manually create free space in the transaction log with a <code>dump transaction</code> command and check on the cause of the failure by referring to the documentation in this manual.
Additional information	Refer to the chapter “Managing Free Space with Thresholds” in the <i>System Administration Guide: Volume 2</i> for information about last-chance thresholds.
Version in which this error is raised	All versions

Error 7414

Severity	10
Message text	Threshold procedure '%.*s' returned an error indication

```
(@status=%d) when called for database '%.*s', segment '%.*s', free_space %ld.
```

Explanation	This error occurs when a threshold procedure is activated but is not able to execute successfully.
Action	<p>Check the Adaptive Server error log for additional information on the reason for the failure. Correct the situation that caused the failure.</p> <p>Execute the failed procedure:</p> <ol style="list-style-type: none">1 Select the database:<pre>1> use <database_name> 2> go</pre><p>Where <i><database_name></i> is the database that contains the failed threshold.</p>2 Determine the stored procedure that is associated with the segment for which the error occurred:<pre>1> sp_helpthreshold <segment_name> 2> go</pre>3 Run the stored procedure manually.
Version in which this error is raised	All versions

Error 7415

Severity	10
Message text	<pre>The transaction log in database %.*s is almost full. Your transaction is being %S_MSG until space is made available in the log.</pre>
Explanation	This error occurs when the last-chance threshold has been exceeded for the specified database. Each transaction accessing the transaction log will receive this message, indicating whether the transaction will be put to sleep or aborted until log space is available.
Action	This message does not ordinarily require action if the transaction was put to sleep and the threshold procedure frees log space. In that circumstance, the sleeping processes will awaken and continue.

If the process was aborted, restart it.

If the process is not able to continue, look in the error log to determine why and manually dump the transaction log.

Version in which this error is raised

All versions

Error 7417

Severity

21

Message text

The last chance threshold entry for logsegment is not present in threshold cache for database '%S_DBID'. Please use 'dbcc thresholds' to see the threshold cache contents and contact your system administrator.

Explanation

Every database in the server has a last-chance threshold, which is an estimate of the number of free log pages required to back up the transaction log. A database's last-chance threshold is stored in the `systhresholds` table, which you can also examine using `lct_admin`.

When you run `dbcc checkcatalog`, and the server determines that the last-chance threshold is missing from memory, and it raises the 7417 error; you may also see error 7411 ("Error 7411" on page 561) and error 7418 ("Error 7418" on page 566).

The server may be unable to find the last-chance threshold under the following circumstances:

- A database has mixed data and log, and
- Certain maintenance activities caused the user table extents to be stored on disk segments that are intended for database log use only (known as stranded extents), and
- The server is shut down so that the next restart triggers a log free space recalculation during recovery.

Due to an Adaptive Server problem, no last-chance threshold is instantiated for the database in this scenario.

Action

Upgrade to the latest version of Adaptive Server Enterprise available for your platform. Updates to Adaptive Server's log free space accounting in the latest versions ensure that the last-chance threshold is correctly handled in the scenario that previously led to the 7417 errors.

Additional information	Stranded extents can result in the server being unable to find the space it expected to find in the transaction log. This can lead to other serious errors, such as error 3475, which causes the server to shut down. See the write-up for error 3475.
Version in which this error is raised	All versions

Error 7418

Severity	22
Message text	The last chance threshold entry for logsegment is not present in systhresholds table for database '%S_DBID'. Please contact your system administrator.
Explanation	<p>Every database in the server has a last-chance threshold. The threshold is an estimate of the number of free log pages required to back up the transaction log. A database's last-chance threshold is stored in the systhresholds table, which you can examine using <code>lct_admin</code>.</p> <p>When you run <code>dbcc checkcatalog</code>, and the server determines that the last-chance threshold is missing from memory, it raises the 7418 error; you may also see error 7411 (“Error 7411” on page 561) and error 7417 (“Error 7417” on page 565).</p> <p>The server may be unable to find the last-chance threshold under the following circumstances:</p> <ul style="list-style-type: none">• A database has mixed data and log, and• Certain maintenance activities caused the user table extents to be stored on disk segments that are intended for database log use only (known as stranded extents), and• The server is shut down so that the next restart triggers a log free space recalculation during recovery. <p>Due to an Adaptive Server problem, no last-chance threshold is instantiated for the database in this scenario.</p>
Action	Upgrade to the latest version of Adaptive Server Enterprise available for your platform. Updates to Adaptive Server’s log free space accounting in the latest versions ensure that the last-chance threshold is correctly handled in the scenario that previously led to the 7418 errors.

Additional information	Stranded extents can result in the server being unable to find space it expected to find in the transaction log. This can lead to other serious errors such as error 3475, which causes the server to shut down. See the write-up for error 3475.
Version in which this error is raised	All versions

Error 7438

Severity	21
Message text	The threshold cache for the database '%.*s' (%d) is corrupt. Please use dbcc dbrepair with option 'newthreshold' to rebuild the threshold cache of the database.
Explanation	Adaptive Server maintains a cache of <code>systhresholds</code> information and threshold pointers for each database. This cache is validated against the <code>systhresholds</code> table when a threshold is added or modified, or when <code>dbcc checkcatalog</code> is run. If the cache validation fails, error 7438 is raised.
Action	Run <code>dbcc dbrepair (dbname, newthreshold)</code> to rebuild the threshold cache. See the <i>Reference Manual: Commands</i> for instructions on using <code>dbcc dbrepair</code> .
Additional information	A configurable shared memory dump (CSMD) configured for error 7438 may be helpful in determining root cause for this error. Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Auditing Errors (7600s)

This section contains error messages pertaining to Adaptive Server Enterprise auditing. Refer to “Auditing” in the *System Administration Guide: Volume 1* for more information about auditing functionality.

Error	Page
Error 7618	569

Error 7618

Severity	20
Message text	The audit process is not currently active. Contact a user with System Administrator (SA) or System Security Officer (SSO) role.
Explanation	<p>When auditing is required, an audit record is allocated on the stack and filled in with information that varies depending on the type of event. Adaptive Server formats the record into an audit buffer (the audit queue) and places this buffer in the audit process mailbox.</p> <p>If the audit process is not active when auditing is required, Adaptive Server terminates the user’s process because the process should have been audited but the audit process was unable to do so.</p>
Action	<p>Resolve the errors that stopped auditing (for example, change the number of locks, or the space in <code>sybsecurity</code>, and so on), then have a user with the Sybase Security Officer (SSO) role restart auditing:</p> <pre>1) sp_audit 'restart' 2) go</pre>

Note If you are running a version of Adaptive Server Enterprise earlier than version 11.9.2.5, you must restart Adaptive Server to restart the audit process.

Additional information	Version 15.0 and later, refer to the “Auditing” chapter in the <i>System Administration Guide: Volume 1</i> for detailed information about auditing; versions earlier than 15.0, refer to the “Auditing” chapter in the <i>System Administration Guide</i> .
Version in which this error is raised	All versions

Remote Procedure Call Errors (8000s)

This section contains error messages for Adaptive Server Enterprise remote procedure calls (RPC).

Error	Page
Error 8006	571
Error 8009	572

Error 8006

Severity	20
Message text	SQL Server error detected by %S_MSG. Refer to the preceding %S_MSG message.
Explanation	<p>Error 8006 is a general message that Adaptive Server can raise in a number of situations when working with another Sybase server. The “%S_MSG” in the message text is the name of the server that generated the error.</p> <p>Most commonly, this error indicates that Backup Server encountered a problem in running the dump database or load database commands. When you issue a dump or load command from Adaptive Server, the server interprets the command and sends remote procedure calls (RPCs) to the Backup Server. As dumps and loads execute, Adaptive Server and Backup Server use RPCs to exchange instructions and status messages. Thus, an 8006 error during a dump or load is typically accompanied by additional messages originating from the Backup Server.</p>
Action	<p>Note the date-time stamp of the 8006 error in the Adaptive Server error log. Check the other server’s error log (for example, the Backup Server’s error log) for error messages corresponding to that time.</p> <p>If there are any Backup Server errors, look up the error description in this manual for assistance. Some common errors are:</p>

- 3.17.2 – Adaptive Server was started from a directory other than its default working directory. See the write-up for “Error 3240” on page 395 for details.
- 4.76.3 – the actual master device size does not match the value in *sysdevices.high*. See the write-up for “Error 5013” on page 476 for a description of how to reset the high value.
- 4.82.2 – there was an error in reading from or writing to a dump device. Check the device and directory permissions.

If no other errors accompanied the 8006 message, there may be a problem with Backup Server. Take these steps:

- 1 Use the `showserver` utility to check that Backup Server is running. Refer to the Adaptive Server utility commands manual for your platform for assistance.
- 2 Log into Adaptive Server and verify that you can communicate with Backup Server, using the following command:

```
1> exec SYB_BACKUP...sp_who
2> go
```

Under normal situations, you should see a few lines of `sp_who` output.

Additional information

See the Adaptive Server Enterprise *Utility Guide* for instructions on using the `showserver` utility.

Version in which this error is raised

All versions.

Error 8009

Severity

16

Message text

Error encountered by %S_MSG. Please refer to %S_MSG messages for details.

Explanation

Error 8009 is a general purpose message that Adaptive Server can raise in a number of situations when working with another Sybase server. The “%S_MSG” in the message text is the name of the server that generated the error. Most commonly, this error indicates that Backup Server encountered a problem in running the dump database or load database commands.

When you issue a `dump` or `load` command from Adaptive Server, the server interprets the command and sends remote procedure calls (RPCs) to the Backup Server. As dumps and loads execute, Adaptive Server and Backup Server use RPCs to exchange instructions and status messages. Therefore, an 8009 error during a `dump` or `load` is typically accompanied by additional messages originating from the Backup Server.

Action

Check the Backup Server and Adaptive Server error logs for additional messages. If there are any Backup Server errors, refer to that error description in this manual for assistance in resolving the error.

If no other errors accompanied the 8009 message, there may be a problem with Backup Server. Take these steps:

- 1 Use the `showserver` utility to check that Backup Server is running. Refer to the Adaptive Server utility commands manual for your platform for assistance.
- 2 Log into Adaptive Server and verify that you can communicate with Backup Server, using the following command:

```
1> exec SYB_BACKUP...sp_who
2> go
```

Under normal situations, you should see a few lines of `sp_who` output.

If Backup Server is down or is not communicating, restart the server and repeat the above steps. If the problem persists, contact Sybase Technical Support.

Additional information

See the Adaptive Server Enterprise *Utility Guide* for instructions on using the `showserver` utility.

Version in which this error is raised

All versions

Descriptor Manager Errors (8200s)

This section contains error messages for the Adaptive Server Enterprise Descriptor Manager.

Error	Page
Error 8201	575
Error 8203	576
Error 8204	577
Error 8207	578
Error 8210	579
Error 8211	580
Error 8219	580
Error 8239	582
Error 8241	582

Error 8201

Severity	26
Message text	Keep count of descriptor (objid=%ld, dbid=%d) was expected to be %d. Instead %d was found.
Explanation	Object descriptors are the data structures used to manage access to Adaptive Server objects. The Descriptor Manager controls and synchronizes access to descriptors. It uses a “keep count” to determine whether a descriptor is in use. This guarantees that a descriptor is protected while in use. Error 8201 occurs if the Descriptor Manager determines that the “keep count” is incorrect.
Action	Shut down and restart Adaptive Server, which should clear error 8201. If the error persists, contact Sybase Technical Support.
Additional information	Have the information listed in “Reporting errors” on page 4 ready if you call Sybase Technical Support:

Version in which this error is raised All versions

Error 8203

Severity 26

Message text Expected to find the descriptor for object %d in database %d in %s state.

Explanation Object descriptors are the data structures used to manage access to Adaptive Server objects. As used, a descriptor passes through a number of different states:

- `free` – the descriptor is on the free resource list. It can be used only by the Descriptor Manager that can change it to an active descriptor.
- `create` – this is the stage of a descriptor between free and active, while the field information is being supplied.
- `active` – the descriptor is accessible from outside the Descriptor Manager. This state does not mean that the descriptor is being used.
- `drop` – an object is in the process of being dropped.
- `destroy` – a descriptor is being given a version so it can be used for another object.
- `clean` – the descriptor is between the active and free states of the chain; in this state the descriptor cannot be found. The next state is either “free” or “create.”

When the Descriptor Manager tries to access a descriptor and the descriptor is in an incorrect state, error 8203 occurs.

Error 8203 occurs with these states values.

Note Some state values may have been added since this book was published. Contact Sybase Technical Support for additional information.

State	Meaning
1	Adaptive Server expected a descriptor it was trying to get to be in state “active,” but it was not.

State	Meaning
2	When Adaptive Server prepared to clean up a descriptor, it expected the descriptor to be in state “destroy,” but it was not.
3	When Adaptive Server prepared to return a descriptor from the resource free list, it expected the descriptor to be in state “free,” but it was not.
4	Adaptive Server prepared to move a descriptor from state “drop” to state “destroy,” but the descriptor was not in state “drop.”
5	Adaptive Server expected a descriptor it was creating for use to be in state “create,” but it was not.
6	Adaptive Server expected a descriptor it was creating for use to be in state “create.” However, Adaptive Server received an invalid status and could not determine what state the descriptor was in. Therefore, even though it expected the descriptor to be in state “create,” the message shows it expected the state to be “Invalid Stat Return.”
7	When Adaptive Server prepared to move a descriptor to state “free,” it expected the descriptor to be in state “clean,” but it was not.
8	When Adaptive Server needed to mark a descriptor as “active” and it had previously marked it for “drop,” it expected the descriptor to be in state “drop” and it was not.
9	When Adaptive Server prepared to move a descriptor from “active” to “drop”, it checked to make sure the descriptor was in state “active,” but it was not.
10	When Adaptive Server prepared to store information about a new descriptor, it expected the descriptor to be in state “create,” but it was not.
11	When Adaptive Server tried to initialize the lock promotion information in the descriptor, it expected the descriptor to be in state “create,” but it was not. This state is only raised by <code>diagserver</code> .

Action	<ul style="list-style-type: none"> If error 8203 occurs with a state other than 6, shut down and restart Adaptive Server to clear the cache corruption. If the 8203 error occurs again, call Sybase Technical Support. If error 8203 occurs with state 6, call Sybase Technical Support.
Additional information	Have the information listed in “Reporting errors” on page 4 ready when you call Sybase Technical Support.
Version in which this error is raised	All versions

Error 8204

Severity	20
Message text	Expected to receive system catalog id. Instead %d was received.

Explanation	<p>Object descriptors are the data structures used to manage access to Adaptive Server objects. Adaptive Server tracks individual descriptors by database table and system catalog ID. Error 8204 occurs if an invalid system catalog ID is returned while Adaptive Server is trying to locate a descriptor.</p> <p>Error 8204 is caused by incorrect system catalogs or an Adaptive Server problem.</p>
Action	Call Sybase Technical Support.
Additional information	When you call Sybase Technical Support, have the information listed in “Reporting errors” on page 4, including the output of <code>dbcc checkcatalog</code> if you know the database in which the error occurred.
Version in which this error is raised	All versions

Error 8207

Severity	17
Message text	<p>Version 12.5 and later</p> <p>Couldn't find an available descriptor for an object. Raise the value of the configuration parameter 'number of open objects' for more descriptors.</p> <p>Version 12.0.x and earlier</p> <p>Couldn't find an available descriptor for an object. Configure your system for more descriptors.</p>
Explanation	<p>This message displays when Adaptive Server is configured for fewer open objects (also known as descriptors) than are currently needed.</p> <p>The configuration parameter <code>number of open objects</code> sets the maximum number of objects that can be open at one time in Adaptive Server.</p>
Action	<p>1 Use the system procedure <code>sp_configure</code> to display the run value of the configuration parameter you need to change:</p> <pre>1> sp_configure "number of open objects" 2> go</pre>

Increasing the value of number of open objects does not have a significant impact on performance or storage requirements. However, be careful when changing Adaptive Server configuration parameters. Refer to the chapter “Setting Configuration Parameters” in the *System Administration Guide: Volume 1* for information about selecting an appropriate value for number of open objects.

- 2 Increase the value of the number of open objects parameter to a higher value than the current run value for that parameter:

```
1> sp_configure "number of open objects", <new_value>
2> go
```

This configuration change is dynamic and takes effect immediately.

Version in which this error is raised

All versions

Error 8210

Severity

21

Message text

Duplicate of work table descriptor found in the id hash table for temporary object %d.

Explanation

Adaptive Server uses work tables to temporarily store query results. Work table descriptors are the data structures used to manage access to work tables. For each task, the server maintains a list of the descriptors for work tables in use by that task. Error 8210 is raised when a problem is detected while placing a descriptor into a task work table descriptor list; the error results in a stack trace.

Error 8210 is raised with these states:

State	Meaning
1	The descriptor was found to belong to an object that is not a work table.
2	There is already a descriptor for this work table in the task's work table descriptor list.

Error 8210 is due to an Adaptive Server problem with batch scripts or stored procedures that use cursors and employ referential integrity checks.

Action

Contact Sybase Technical Support.

Additional information Have the information listed in “Reporting errors” on page 4 when you call Sybase Technical Support.

Note The error log contents are especially critical for analysis by Sybase Technical Support. Additional messages that may appear in the error log, along with the 8210 message, may be important clues to the underlying problem.

Version in which this error is raised All versions.

Error 8211

Severity 26

Message text Mismatch found between the name and id descriptor hash table. Descriptor hashed by name = 0x%0x and hashed by id = 0x%0x.

Explanation Object descriptors are the data structures used to manage access to Adaptive Server objects. ID and name hash tables are used to maintain object descriptors. Every object descriptor should have a corresponding entry in the ID and name hash tables. Error 8211 occurs when Adaptive Server finds two descriptors for one object or when a descriptor for an object exists in only one of the hash tables (ID or name table).

Action Shut down and restart Adaptive Server to clear the cache corruption.
If the 8211 error persists, contact Sybase Technical Support.

Additional information Have the information listed in “Reporting errors” on page 4 when you call Sybase Technical Support.

Version in which this error is raised All versions

Error 8219

Severity 18

Message text Couldn't compute lock promotion level in DES for object %d in database %d. Using default level = %d.

Explanation	<p>Use the following configuration parameters to control how Adaptive Server handles lock promotion:</p> <ul style="list-style-type: none"> • lock promotion hwm – sets the maximum number of locks allowed on a table. When the number of locks that an individual command acquires exceeds this number, Adaptive Server attempts to acquire a table lock. • lock promotion lwm – sets the minimum number of locks that are allowed on a table before Adaptive Server attempts to acquire a table lock. Adaptive Server does not attempt to acquire a table lock until the number of locks on a table is equal to the lock promotion lwm. • lock promotion pct – sets the percentage of locks (based on the table size) above which Adaptive Server attempts to acquire a table lock when the number of locks is between the lock promotion hwm and the lock promotion lwm. <p>If a command requires more locks than the number specified by lock promotion lwm but less than the number specified by lock promotion hwm, Adaptive Server compares the number of locks to the percentage of locks on the table. If the number is greater than the percentage, Adaptive Server attempts to issue a table lock.</p> <p>Adaptive Server computes the lock promotion threshold as a percentage of the table size. If Adaptive Server is unable to get the table size, error 8219, an informational message, is displayed and Adaptive Server uses the value of lock promotion lwm to determine the lock promotion level.</p>
Action	<p>Error 8219 can occur if the table size is 0 (zero rows) or the number of pages for the table is 0 (zero). Run dbcc checktable and dbcc checkalloc for the table to determine if this is the case.</p> <p>If this is not the case and the 8219 error occurs again, contact Sybase Technical Support.</p>
Additional information	<p>Refer to “Lock promotions” in the chapter “Monitoring Performance with sp_sysmon” in <i>Performance and Tuning: Monitoring and Analyzing</i>.</p> <p>Have the information listed in “Reporting errors” on page 4 when you call Sybase Technical Support, including the output of sp_configure (or a copy of the configuration file), sp_helpdb, sp_help on the table involved, and sp_spaceused on the table involved.</p>
Version in which this error is raised	<p>Versions earlier than 12.5.x</p>

Error 8239

Severity	20
Message text	Not enough memory to create the cache binding information for `%.*s` in database `%.*s`.
Explanation	A failure occurred when trying to allocate memory from the data cache fragment pool.
Action	The size of the data cache fragment pool is calculated at startup time based on the configuration parameter <code>number of open objects</code> . Increasing that value and restarting Adaptive Server may help avoid this error.
Additional information	A configurable shared memory dump (CSMD) configured for error 8239 may be helpful in determining root cause for this error. Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.5.2 and later

Error 8241

Severity	21
Message text	Fatal internal error encountered while instantiating DBTABLE for database `%.*s`. \'
Explanation	The server encountered an error attempting to clean out and remove all system and user descriptors linked to the <code>dbtable</code> for the given database. The <code>dbtable</code> structure storing information for each open database.
Action	<p>If this error occurs during a database load, it may have been caused by a corrupted dump. Make sure that the versions being used of Adaptive Server and Backup Server are compatible, and retry the dump and load process.</p> <p>If this error occurs during normal operations, verify that the number of open databases and the number of open objects are configured with an appropriately high value.</p>
Additional information	A configurable shared memory dump (CSMD) configured for error 8241 may be helpful in determining root cause for this error. Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.

Version in which this
error is raised

Version 15.0 and later

sysindexes Manager Errors (8400s)

This section contains write-ups for `sysindexes` Manager errors.

Error	Page
Error 8402	585
Error 8419	587

Note Adaptive Server Enterprise version 15.0 (and later) still uses the `sysindexes` sytem table, however some of the information that was stored in `sysindexes` prior to version 15.0 is now stored in the new 15.0 `syspartitions` system table. See the chapter “System Tables” in the Adaptive Server Enterprise version 15.0 *Reference Manual: Tables* guide for updated information about `sysindexes` and `syspartitions`.

Error 8402

Severity	21
Message text	Index row for object %ld (index id %d) was not found in database %d.
Explanation	<p>Each database has a <code>sysindexes</code> table. The <code>sysindexes</code> Manager maintains and provides access to the rows of the <code>sysindexes</code> table. It handles inserting and deleting of <code>sysindexes</code> rows and manages the abort and recovery of individual operations performed on the rows. It also provides read/write access to the database information structure.</p> <p><code>sysindexes</code> contains one row for each:</p> <ul style="list-style-type: none"> • Clustered index • Nonclustered index • Table that has no clustered index

- Table that contains text or image columns

Error 8402 occurs when one of these rows is expected but not found for the object specified in the error message.

Action

- 1 Determine the name of the database mentioned in the error message:

```
1> use master
2> go
```

```
1> select db_name(<database_ID>)
2> go
```

Where *<database_ID>* is the database ID in the error message.

- 2 Search the *sysindexes* table for the object displayed in the error message:

```
1> use <database_name>
2> go

1> select * from sysindexes
2> where id = <object_ID>
3> go
```

Where *<object_ID>* is the object ID in the error message.

- 3 If the object is found in *sysindexes*, continue to step 4.
If the object is not found in *sysindexes*, call Sybase Technical Support.
- 4 If the object specified in the error message is not a system table (a system table's object ID is less than 100), continue to step 5.
If the object specified in the error message is a system table, refer to "Fixing a corrupted system table index" in the chapter "Other Useful Tasks" in the most recent version of the *Troubleshooting and Disaster Recovery* guide for instructions on how to repair the system table index, then go to step 7.
- 5 Drop the index for the object specified in the error message.
- 6 Re-create the index. This clears the corruption in most cases.
- 7 Run *dbcc checktable* on the table to verify that the corruption is gone. If corruption still exists, call Sybase Technical Support.

Additional information

Refer to "drop index" and "create index" in the *Reference Manual: Commands* for information about dropping and re-creating indexes.

Version in which this error is raised

All versions

Error 8419

Severity 21
 Message text Could not find index descriptor for objid %d, indid %d in dbid %d.

Explanation Each index accessed in Adaptive Server requires a descriptor that is a “handle” for referencing the index. When the descriptor is no longer needed (for example, the query completes or the index is dropped), the server frees up the descriptor so that it can be assigned to another object as needed.

Error 8419 prevents a stored procedure or other compiled object from being recompiled after an index or a table column is dropped; consequently the server attempts to process the query via a stale plan, using the missing index or column. When the descriptor containing *sysindexes* information for the index is found to be missing from cache, error 8419 is raised. Here is a typical scenario:

```
1> create table t1
2> create procedure p1 as insert into t1 values(...)
3> go

1> begin transaction
2> go

1> create clustered index on t1(..)
2> execute procedure p1 (query plan uses the new index)
3> rollback transaction (index dropped)
4) go

1> execute procedure p1 (Error 8419 raised as the server
                        attempts to use the old plan)
```

Error 8419 is due to an Adaptive Server problem and is raised with the following states.

Note All states do not occur in all versions of the product and additional state may have been added or removed since this book was published. Contact Sybase Technical Support for more information.

State	Meaning
1	Index descriptor row not found in cache when attempting to make a copy of the table's first index structure.
2	Index descriptor row not found in cache when attempting to copy the next index structure for the table.

State	Meaning
3	Could not retrieve a pointer to the first index descriptor row for the table in the cache.
4	Could not retrieve a pointer to the next index for the table in the cache as expected.
5	Could not retrieve a specific field within an index descriptor row using a pointer to the table's first index row.
6	Could not retrieve a specific field within an index descriptor row using a pointer to the table's next index row.
7	Could not retrieve a pointer to an index descriptor row in the cache after dropping the index.
8	Could not retrieve a pointer to an index descriptor row in the cache after dropping a column of the table.
9	Could not retrieve a pointer to an index descriptor row in the cache while accessing the root page of the index. This state is raised only by <code>diagserver</code> .
10	Could not retrieve a pointer to an index descriptor row in the cache while releasing locks associated with the row.
11	Could not find the start of the index descriptor chain.
12	Could not update a specific field within an index descriptor row using a pointer to the index row.

Action	Drop and recompile the stored procedure hitting the error. If that does not resolve the problem, contact Sybase Technical Support.
Additional information	Have the information listed in "Reporting errors" on page 4 when you call Sybase Technical Support.
Version in which this error is raised	All versions

Process Status Structure Errors (8600s)

This section contains error messages for the Adaptive Server Enterprise Process Status Structure (PSS).

Error	Page
Error 8601	589

Error 8601

Severity	20
Message text	Process %d tried to keep the PSS for process %d when it already had the PSS for process %d kept.
Explanation	<p>A process is a task being carried out by Adaptive Server. Processes can be in various states, such as running or waiting for resources. Processes cease to exist when they complete normally or are terminated with the <code>kill</code> command. For each user process and some system processes, Adaptive Server maintains a Process Status Structure (PSS) to track the state of the process and its use of resources.</p> <p>Typically, each process is concerned with its own PSS. However tasks such as diagnostic commands, system procedures like <code>sp_who</code>, and the <code>kill</code> command require that a process interact with the PSS for another process. For example a process performing a <code>kill</code> must temporarily attach (“keep”) the PSS of the target process to its own PSS, read or write necessary information in the target PSS, and “unkeep” it. A process can only keep one PSS (besides its own) at a time.</p> <p>Error 8601 is raised when a process is already keeping the PSS of one process, and tried to keep the PSS of another process at the same time. The error is sometimes seen when the System Administrator or other user is attempting a <code>kill</code> operation or running a diagnostic command.</p>
Action	To correct this error:

- Shut down and restart Adaptive Server to clear the PSS structure.
- Defer kill operations and diagnostic commands to a later time.

Version in which this
error is raised

All versions

Site Buffer Manager Errors (8700s)

This section contains error messages for the Adaptive Server Enterprise Site Buffer Manager.

Error	Page
Error 8704	591

Error 8704

Severity	26
Message text	Could not get the address lock. address=0x%1x, status=%d.
Explanation	<p>Address locks control memory addresses. They are used for:</p> <ul style="list-style-type: none"> • Object allocation map (OAM) pages • Allocation pages • Updating the disk maps in the database table memory structure • Updating values for descriptors and site buffers <p>The Site Buffer Manager is one of the managers that control Server-to-Server communications.</p> <p>Error 8704 occurs when the Site Buffer Manager fails to get an address lock for an active site buffer and fails to attach a logical connection to it. This error is caused by an Adaptive Server problem.</p>
Action	<p>Shut down and restart Adaptive Server.</p> <p>If this error persists, setting up a configurable shared memory dump (CSMD) may be helpful in determining the root cause for this error. Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of <i>Troubleshooting and Disaster Recovery</i>.</p>

Error 8704

Additional information	Have the information available listed in “Reporting errors” on page 4, including the output of <code>sp_configure</code> , <code>sp_who</code> , and <code>dbcc lock</code> before calling Sybase Technical Support.
Version in which this error is raised	All versions

This section contains write-ups of common Adaptive Server Enterprise Disk Manager errors.

Error	Page
Error 9004	593

Error 9004

Severity	20
Message text	I/O error detected on disk buffer 0x%X for page request %ld, dbid %d.
Explanation	When you issue dbcc commands such as dbcc checkalloc or dbcc checkdb, Adaptive Server reads a specified number of pages from disk into disk buffers so that consistency checks can be carried out on the pages. Error 9004 is raised when the server attempts a synchronous disk read during dbcc processing, but the request fails with an I/O error.
Action	Check the server error log for similar errors from other operations, such as an 823 error when running a query ("I/O error detected during %S_MSG for %S_BUF.") or sddone errors ("sddone:write error on virtual disk %d"). Refer to the write-ups for those errors elsewhere in this manual for more information. Since Adaptive Server reports 9004 and similar errors when its I/O operation receives an unsuccessful return value, these errors can be an early warning of disk problems.

Check your disks as follows:

- Note the database ID shown in the 9004 message text.
- Identify the affected devices using the steps shown in “Determining the physical device on which a database resides” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.

- Examine the availability and condition of the devices using your standard operating system procedures.

Additional information

See “Checking the operating system error log” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* if you need assistance in locating the system log.

Version in which this error is raised

All versions

Replication Agent Manager Errors (9100s - 9200s)

This section contains write-ups for common errors that can occur when Adaptive Server Enterprise is used with Replication Server®.

Error	Page
Error 9122	595
Error 9287	596
Error 9289	597
Error 9290	599

Error 9122

Severity	21
Message text	The log transfer process lost its connection to the LTM for database '%.*s'. Once the LTM re-establishes the connection, replication will continue with no data loss.
Explanation	<p>Through version 11.0.3 of Replication Server, the Log Transfer Manager (LTM), a replication system component, reads the transaction log of the Adaptive Server database being replicated and transfers this information to the Replication Server.</p> <p>Error 9122 is raised if Replication Server or the LTM are shut down improperly; if the shut down occurs while sending data, Adaptive Server is unable to disconnect the LTM connection normally, resulting in a 9122 error and stack trace.</p>
Action	<p>Error 9122 is not a serious error. Once you restart Replication Server and the LTM, the log transfer process resumes normally.</p> <p>To prevent future occurrences of this error, perform a log suspend before shutting down the LTM.</p>

Version in which this error is raised 11.0.3 and earlier

Error 9287

Severity 19

Message text %S_REPAGNT. The data at (%d, %d) is skipped because it contains data wider than the Replication Server, '%.*s', can handle. Check the 'data limits filter mode' configuration parameter..

Explanation Error 9287 is raised if:

- Replication Server does not support large identifiers
- The `sp_config_rep_agent` configuration parameter `data limits filter mode` is set to “`skip`”, and
- One or more columns are skipped during replication

Refer to Appendix B in the *Replication Server Design Guide* for more information about Replication Server and log transfer language (LTL) compatibility.

Action Upgrade Replication Server to a version where the LTL version is 700 or later to support the larger identifiers and column size, or avoid replicating tables that use wide columns and large identifiers.

Additional information The `data limits filter mode` Replication Agent configuration parameter specifies how Replication Agent handles log records containing new, wider columns and parameters, or larger column and parameter counts, before attempting to send them to Replication Server. Valid values are:

- `off` – Replication Agent allows all log records to pass through.
- `stop` – Replication Agent shuts down if it encounters log records containing wide data.
- `skip` – Replication Agent skips log records containing wide data and prints a message to the error log.

- `truncate` – Replication Agent truncates wide data to the maximum length that the Replication Server can handle.

Warning! Warning: Sybase recommends that you do not use the `data limits filter mode "off"` setting with Replication Server 12.1 or earlier, which can cause Replication Agent to skip or truncate the wide data or to stop.

The default value of `data limits filter mode` depends on the Replication Server version. For Replication Server versions 12.1 and earlier, the default value is "stop." For Replication Server versions 12.5 and later, the default is set to "off."

Version in which this error is raised

Version 12.5 and later

Error 9289

Severity	20
Message text	<code>%S_REPAGNT. Log record (%d, %d) for object = %d contains inconsistent data. Transaction log may be corrupt. Please contact Sybase Technical Support.</code>
Explanation	While scanning the log, Replication Agent found an invalid column or parameter length. Replication Agent for this database is shut down.
Action	Perform a full database dump as soon as possible. Because this is a permanent error in the log, Replication Agent will shut down again if it is restarted. This problem could affect one or more log records, but the transaction log is not compromised.

Take the following actions:

- 1 Enable trace flag 9217:

```
use <db_name>
sp_config_rep_agent <db_name>, 'traceon', '9217'
sp_start_rep_agent <db_name>
```

By default, Replication Agent stops replication after reporting error 9289. When you set trace flag 9217, Replication Agent attempts to replicate data included in the inconsistent log record. If the replication continues, the data being replicated may be incorrect. When processing the affected log record, Replication Agent sends message 9290:

```
RepAgent(%d). Log record (%d, %d) for object =
xxxxxxx contains inconsistent data. Rep Agent will
try to extract the information read and continue
replication. Unexpected data may be replicated.
```

This trace is specifically for systems that require nonstop replication, even if data consistency is not maintained.

Once Replication Agent reports error 9290:

- Identify the object name using the object ID indicated in the error.
- Use `rs_subcmp` or any other customized tool to compare data between primary and replicated sites.

All affected objects must be synchronized using `rs_subcmp`.

2 Reset the truncation point.

If the trace flag 9217 does not solve the problem, reset the secondary truncation point to skip the log records affected by this problem because Replication Agent is down due to an error.

Use the following procedure to set the secondary truncation point to the end of the transaction log:

a Enter:

```
use <db_dbname>
dbcc settrunc(ltm,ignore)
dbcc settrunc(ltm,'end')
```

b Shut down Replication Server.

c Connect to the Replication Server System Database (RSSD) and execute:

```
rs_zeroltm <dsname>, <dbname>
```

d Restart Replication Server.

e Start Replication Agent.

Note This procedure requires a full synchronization of the database.

Additional information

If it is necessary to troubleshoot further, perform a dump of the transaction log containing the problem, which allows Sybase Technical Support to further analyze the issue.

Version in which this error is raised

Version 12.5.3 and later

Error 9290

Severity	18
Message text	%S_REPAGNT. Log record (%d,%d) for object = %d contains inconsistent data. Rep Agent will try to extract the information read and continue replication. Unexpected data may be replicated.
Explanation	<p>Error 9290 is raised only when Replication Agent trace flag 9217 has been activated. While scanning the log, Replication Agent has found an invalid column or parameter length stored in the log record.</p> <p>This error is similar to error 9289, however when this trace flag is active, Replication Agent attempts to replicate the information stored in the inconsistent log record.</p>
Action	The transaction, at the indicated log record, may have replicated inconsistent data. Please contact Sybase Technical Support for instructions on running the dbcc log command to identify the transaction.
Additional information	If it is necessary to troubleshoot further, take a dump of the transaction log containing the problem to allow Sybase Technical Support to further analyze the issue.
Version in which this error is raised	Version 12.5.3 and later

Component Integration Services Errors (11200s)

This section contains error messages generated by Component Integration Services (CIS).

Error	Page
Error 11202	602
Error 11203	602

Error 11201

Severity	16
Message text	A remote table definition must be defined before the 'existing' option is used (object '%.*s').
Explanation	The create existing table statement was used without specifying a location clause.
Action	Use the location clause, or remove the “existing” keyword from the create table command: <pre>1 create existing table tablename (column_list) at 'server.database.user.object '</pre> <pre>2 create table tablename (column_list) other_options</pre>
Version in which this error is raised	All versions

Error 11202

Severity	16
Message text	Unable to create table '%.*s' because it has no

remote storage location defined. Creation of local user tables is allowed only in the master and tempdb databases.

Explanation

Component Integration Services (CIS) is a feature that extends Adaptive Server capabilities and allows you to access and manipulate tables in remote servers as if they were local tables. When you install Adaptive Server, files needed to interact with CIS are placed in the installation directory.

Error 11202 is raised when you attempt to create a table, but Adaptive Server is unable to locate or otherwise access one of these files. The problem can occur even if you are creating a local table and do not utilize CIS.

Action

Use the following checklist to ensure that the necessary file is accessible:

- 1 Check that the *SYBASE* environment variable is properly set. If you recently upgraded, make sure that this variable was updated. Restart the server after checking this variable.
- 2 Make sure that you start the server as user “sybase”.
- 3 Check that the files *dsopt1* (necessary for local tables) and *dsopt2* (necessary for remote tables) exist in your installation directory. For example, in UNIX sites the file *\$SYBASE/install/dsopt1* must exist with read permissions before you can create local tables. Restore any missing files and restart Adaptive Server.

Note The Adaptive Server Enterprise directory structure changed with version 12.0. Check the installation and configuration guides for your platform for details.

Version in which this error is raised

All versions

Error 11203

Severity

16

Message text

The command '%s' is not legal for objects of this type.
External definition: %.*

Explanation	<p>Adaptive Server 12.5 enables SQL access to file system directories and the underlying files. A Component Integration Services (CIS) proxy table, created with the following command syntax maps column names to attributes of files that exist within the external directory.</p> <pre style="margin-left: 40px;">create table table-name external file at "directory-path" ...</pre> <p>Given the nature and purpose of CIS proxy tables defining external file systems, not all Adaptive Server commands and functions are applicable to these tables. Error 11203 is raised when the server is instructed to perform an operation that is incompatible with this type of table.</p> <p>For example, error 11203 may be raised when one of the following operations is attempted on an external file system proxy table:</p> <ul style="list-style-type: none"> • dbcc checks • alter table command • delete table command • create index command • Partitioning commands <p>Error 11203 is also raised when you attempt to run an update statistics command on a proxy table mapped to a remote procedure.</p>
Action	<p>Contact Sybase Technical Support. Do not attempt the operation that raised error 11203. Batches or scripts that perform mass operations such as dbcc checks may need to be altered to ensure that the offending operation is not attempted on the external file system proxy table.</p>
Additional information	<p>Have the information listed in “Reporting errors” on page 4 when you call Sybase Technical Support.</p>
Version in which this error is raised	<p>Version 12.5 and later</p>

Configurable Shared Memory Dump Manager (12000s, 12300s)

The configurable shared memory dump (CSMD) tool allows you to set up conditions for capturing shared memory dumps when needed. This chapter describes errors that can occur as a result of performing a CSMD.

Error	Page
Error 12036	605
Error 12326	606
Error 12328	607
Error 12329	607
Error 12330	608
Error 12333	609
Error 12334	609
Error 12335	610

Error 12036

Severity

17

Message text

Collection of monitoring data for table '%.*s' requires that the %s configuration option(s) be enabled. To set the necessary configuration, contact a user who has the System Administrator (SA) role.

Explanation

By default, Adaptive Server does not collect the monitoring information required by the monitoring tables. You must use `sp_configure` to configure Adaptive Server to start collecting monitoring information (`sp_configure enable monitoring`). The monitoring configuration parameters that are required display when you enter:

```
sp_configure Monitoring
```

	When a monitoring table is accessed, Adaptive Server checks to see if the appropriate configuration parameters for that table are enabled. If they are not, Adaptive Server raises error 12036.
Action	A user with the “sa” role must use <code>sp_configure</code> to enable monitoring in Adaptive Server tables and set the appropriate configuration parameters. If the problem persists, set up a configurable shared memory dump (CSMD), which may be helpful in determining root cause for this error message. See the section “Obtaining a CSMD” in the <i>Troubleshooting and Disaster Recovery</i> guide for instruction in setting up configurable shared memory dump.
Additional information	Refer to “Adaptive Server configuration options” in the chapter “Monitoring tables,” in <i>Performance and Tuning: Monitoring and Analyzing</i> for more information.
Version in which this error is raised	Versions 12.5.0.3 and later

Error 12326

Severity	21
Message text	Table Corrupt: Free offset in page header is not reasonable; free offset should be >= %1! and <= %2!; check this page and offset (page#=%3! freeoffset on the page header=%4!).
Explanation	While attempting to delete a row in an all-pages-locked (APL) table the value in the page header for the free offset was found to be outside the legal values.
Action	To get the header of the page with the error do the following: <ol style="list-style-type: none">1 Use the <code>dbcc page</code> command to display information about the page in question. Here is the syntax:<pre>dbcc traceon (3604) dbcc page (<database_name>, <page_number>)</pre>2 Check the HEADER section for the object ID. Adaptive Server Enterprise version 15.0 and later displays the object ID in the <code>ptnid</code> field. Adaptive Server Enterprise versions earlier than 15.0 display the object ID in the <code>objid</code> field. Once you know the object ID, get the name of the object in the affected database:<pre>select object_name(objectid)</pre>

	3	Run <code>dbcc checktable</code> on that object to check for any corruption.
Additional information		This error may be the result of a hardware error or misconfigured database device. Depending on the extent of corruption, you may need to recover from backup or you may be able to rescue the data. See the section “Rescuing data from a corrupted table” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised		Version 12.5.1 and later

Error 12328

Severity	21
Message text	The calculated row offset <code>`%d`</code> of an index key on page <code>`%ld`</code> for object <code>`%ld`</code> in database <code>`%d`</code> is out of bounds (table offset on page = <code>`%d`</code> , page status = <code>`%d`</code> , local offset table address = <code>0x%lx</code>).
Explanation	A row offset (that is, the position on the page that a row starts at) has been calculated; but it is either too small or too large to fit on the page in question.
Action	This is most likely a transient error caused by two jobs trying to build an offset table at the same time. Try re-running the failed query. If the problem persists, set up a configurable shared memory dump (CSMD), which may be helpful in determining root cause for this error message.
Additional information	Setting trace flag 604 will print out more information about the state of the index search being done. See the section “Obtaining a CSMD” in the <i>Troubleshooting and Disaster Recovery</i> guide for instruction in setting up configurable shared memory dump.
Version in which this error is raised	Version 12.5.2 and later

Error 12329

Severity	20
Message text	An attempt was made to fetch logical page <code>`%d`</code> for

`%S_PTNINFO`. However, page belongs to a different partition (`indid`%d` partitionid`%d``) of the same object. This indicates an ASE internal error and does not indicate data corruption.

Explanation	While checking the validity of a page, Adaptive Server found a discrepancy on the partition ID. When the partition ID from the page matches an existing partition ID for the same table, error 12329 is raised. It does not indicate a data corruption, but indicates some Adaptive Server Enterprise internal error.
Action	The action you take depends on the value of <code>State</code> in the error message: <ul style="list-style-type: none">• State 1 – the error occurred during a read of a page from disk. This may indicate an incorrect partition id on the page. If the error persists it may be necessary to unpartition and repartition the table.• State 2 – the error occurred during an attempt to handle a buffer in memory. Re-run the affected query.
Additional information	If the 12329 error persists, it may be necessary to gather a configurable shared memory dump (CSMD). See the section “Obtaining a CSMD” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 12330

Severity	20
Message text	An attempt was made to fetch logical page <code>`%d`</code> with partition ID <code>`%d`</code> , index ID <code>`%d`</code> from <code>%S_PTNINFO</code> . No partition with ID <code>`%d`</code> was found in this object. This indicates an ASE internal error and does not indicate data corruption.
Explanation	The partition mismatch occurred during an attempt to handle a buffer in memory. This indicates a fatal Adaptive Server internal error (and not data corruption on disk.)
Action	It is possible that the 12330 error indicates a transient problem in Adaptive Server. Re-run the affected query.
Additional information	If error 12330 persists, it may be necessary to gather a configurable shared memory dump (CSMD). See “Obtaining a CSMD” in the <i>Troubleshooting and Disaster Recovery</i> guide.

Version in which this error is raised	Version 15.0 and later
---------------------------------------	------------------------

Error 12333

Severity	20
Message text	A global index '%1!' is in suspect state. Please drop and re-create the index.
Explanation	This message indicates that a global index has been accessed which is marked as suspect. For example, it may have been created with a sort order that is different than the server's default sort order.
Action	Drop and recreate the indicated index.
Additional information	If the 12333 error persists after recreating the index contact Sybase Technical Support. Have the definition of the index and its associated table available. It may be necessary to gather a configurable shared memory dump (CSMD). See “Obtaining a CSMD” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 12334

Severity	21
Message text	Attempts to read logical page `%ld`, virtpage `%ld` from virtual device `%d` for %S_PTNINFO failed due to read failure in prefetch. Please contact Sybase Technical Support.
Explanation	A problem occurred when another process tried to read the indicated page asynchronously. Adaptive Server will retry this operation 10 times, then raise a 12334 error.
Action	It may be possible that some other process is taking an excessive amount of time while accessing the indicated page. Check either the <code>sysprocesses.cpu</code> field. For example: <pre>select top 20 spid, status, program_name, cpu from master..sysprocesses</pre>

```
order by cpu desc
```

Or, if you have Monitoring and Diagnostic (MDA) tables installed, check the `monProcessActivity.CPUTime` field to check for jobs with high CPU usage. For example:

```
select top 20 a.SPID, a.CPUTime, b.Application, b.Command
  from master..monProcessActivity a, master..monProcess b
  where a.SPID = b.SPID
order by a.CPUTime desc
```

Examine any jobs with unexpectedly high cpu usage to see what they are doing. When any problems are corrected, re-run the affected query.

Additional information

If error 12333 persists, it may be necessary to gather a configurable shared memory dump (CSMD). See “Obtaining a CSMD” in the *Troubleshooting and Disaster Recovery* guide section “Other Useful Tasks.”

Version in which this error is raised

Version 15.0 and later

Error 12335

Severity

21

Message text

```
An attempt was made to read logical page '%d' for object '%d' in database '%d'. The page was not read successfully
```

```
because of repeated failures to acquire the page using asynchronous prefetch.
```

Explanation

Explanation

A problem occurred when another process tried to read the indicated page asynchronously. Adaptive Server checks whether the retry threshold (10 times) for the page has been exceeded and raises error 12335 if it has.

Action

It is possible that this condition is due to an Adaptive Server resource shortage. Ensure that the configuration parameters `max async i/os per server`, `max async i/os per engine`, and `disk i/o structures` are set high enough for the volume of I/O is being done.

Additional information	In Adaptive Server 15.0 and later, error 12334 is raised instead of 12335. See “Error 12334” on page 609. If the 12335 error persists, it may be necessary to gather a configurable shared memory dump (CSMD). See “Obtaining a CSMD” in the <i>Troubleshooting and Disaster Recovery</i> guide section Other Useful Tasks.”
Version in which this error is raised	Version 12.5.4 (pre-15.0)

Fault Isolation Utility Errors (12700s)

This section contains error messages for the Adaptive Server Enterprise recovery fault isolation (RFI) functionality.

Error	Page
Error 12716	613
Error 12717	614

Error 12716

Severity	17
Message text	Data unavailable: Unable to complete because the required data is currently unavailable. The command failed and the transaction aborted. Try again later or contact your System Administrator.
Explanation	<p>The recovery fault isolation (RFI) feature enables recovery to mark a single page or a set of pages of a user database suspect, while allowing access to the remainder of the database. The bad pages are said to be “offline,” and may not be used until the problem is corrected.</p> <p>Error 12716 is raised when you issue a query that requires access to an offline page. This error is followed by error 12717 (Suspect data encountered: Database '%.*s', object '%.*s', index ID '%ld', page ID '%ld').</p>
Action	<p>The offline pages must be repaired before they can be used. Responsibility for repairing offline pages belongs to a user with <code>sa_role</code>. The System Administrator can bring a page online to make it accessible to <code>sa_role</code>:</p> <pre>1> sp_forceonline_page <database_name>, <pgid>, 'sa_on' 2> go</pre>

Where *<pgid>* is the page ID shown in the accompanying 12717 error.

Warning! When forcing pages online, Adaptive Server does not perform any consistency checks before the pages are forced online.

Use existing methods to diagnose and repair problems, checking the server error log for any errors related to the offline page and determining appropriate corrective action. Before making repaired pages available to all users, check that repairs have been made correctly by running `dbcc` commands or queries with known results.

Additional information	For a detailed discussion of RFI, refer to “Online recovery and fault isolation” in the chapter “Server Recovery Tasks” in the most recent version of <i>Troubleshooting and Disaster Recovery</i> .
Version in which this error is raised	All versions

Error 12717

Severity	17
Message text	Suspect data encountered: Database '%.*s', object '%.*s', index ID '%ld', page ID '%ld'
Explanation	<p>Adaptive Server’s recovery fault isolation (RFI) feature enables recovery to mark a single page or a set of pages of a user database suspect, while allowing access to the remainder of the database. The bad pages are said to be “offline”, and may not be used until the problem is corrected.</p> <p>Error 12717 is raised when you issue a query that requires access to an offline page. This error is preceded by error 12716 (see “Error 12716” on page 613).</p>
Action	<p>The offline pages must be repaired before they can be used. Responsibility for repairing offline pages belongs to a user with <code>sa_role</code>. The System Administrator can bring a page online to make it accessible to <code>sa_role</code>:</p> <pre>1> sp_forceonline_page <database_name>, <pgid>, 'sa_on' 2> go</pre>

Where *<pgid>* is the page ID shown in the 12717 error.

Warning! When forcing pages online, note that Adaptive Server does not do any consistency checks before putting pages online.

Use existing methods to diagnose and repair problems, checking the server error log for any errors related to the offline page and determining appropriate corrective action. Before making repaired pages available to all users, check that repairs have been made correctly by running *dbcc* commands or queries with known results.

Additional information

For a detailed discussion of RFI, refer to “Online recovery and fault isolation” in the chapter “Server Recovery Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.

Version in which this error is raised

All versions

Partition (14100s) and Partition Key Manager (14300s) Errors

Partitions, like a large-scale index, provide faster and easier access to data. In particular, partitioning can help manage large tables and indexes by dividing them into smaller, more manageable pieces.

This chapter covers resolutions for errors that can occur when using partitions in Adaptive Server. Refer to the chapter “Partitioning Tables and Indexes” in the *Transact-SQL User’s Guide* for more information about using partitions.

Error	Page
Error 14100	617
Error 14108	618
Error 14120	618
Error 14124	619
Error 14126	620
Error 14313	620

Error 14100

Severity	21
Message text	Partition row for database %s, object %s, index %s, partition %s not found.
Explanation	Adaptive Server failed to locate the partition row.
Action	Error 14100 indicates an internal Adaptive Server Enterprise issue, which may be caused by an insufficient number of open partitions. Increase the configuration for number of open partitions and retry the affected query. See “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for details on setting this parameter. If the problem persists, contact Sybase Technical Support.

Additional information	If you contact Sybase Technical Support, be prepared to provide the stack trace from the error log, as well as any information about what was occurring with the spid that hit the error. If the 14100 error persists, it may be necessary to gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 14108

Severity	21
Message text	Could not find partition descriptor for objid %d, indid %d, ptnid %d in dbid %d.
Explanation	Partition descriptor is an Adaptive Server Enterprise internal data structure used to keep track of partition state, including the <code>syspartitions</code> row for a given partition. This is an Adaptive Server internal error raised when a database task fails to grab the partition descriptor used to set up a table.
Action	Error 14108 may be due to an insufficient number of open partitions. Increase the configuration for number of open partitions and retry the affected query. See “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for details on setting this parameter.
Additional information	If the 14108 error persists, it may be necessary to gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 14120

Severity	26
Message text	Expected to find the partition descriptor for objid `%d` indid `%d` ptnid `%d` in dbid `%d`, in a `%s` state

(found state `~%d` instead).

Explanation	Before executing certain operations on partitions, Adaptive Server Enterprise checks the state of the partition descriptor (PDES) to verify that it is in a valid state for the desired operation. Error 14120 is raised when the state that is found is not valid or what was expected.
Action	Error 14120 indicates an internal Adaptive Server Enterprise issue. Depending on the table involved, it may be necessary to reboot Adaptive Server to clear the invalid state. If <code>sp_monitorconfig</code> number of open objects shows a <code>Num_reuse count > 0</code> , increasing the number of open objects parameter may lessen the chances of error 14120 re-occurring.
Additional information	Contact Sybase Technical Support. If the error persists, it may be necessary to gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of the <i>Adaptive Server Enterprise Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 14124

Severity	20
Message text	The to be installed PDES existed already. Dbid %1!, objid %2!, partition id %3!, state %4!.
Explanation	After inserting a new row into the <code>syspartitions</code> system table, the corresponding partition descriptor (PDES) is instantiated. Error 14124 indicates that the PDES already exists in memory.
Action	<ul style="list-style-type: none"> • Use <code>dbcc checktable</code> to check the indicated object. If you find inconsistencies, correct them first. See the section “SQL Reference, dbcc commands” in the <i>Component Integration Services User’s Guide</i>. • If the PDES in memory is stale information that remains by mistake, reboot Adaptive Server can help to solve the problem. • Error 14124 can also indicate an internal Adaptive Server Enterprise error. Contact Sybase Technical Support

Additional information	If error 14124 persists, it may be necessary to gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 14126

Severity	17
Message text	Couldn't find an available partition descriptor. Raise the value of the configuration parameter 'number of open partitions' for more partition descriptors.
Explanation	In Adaptive Server Enterprise version 15.0 and later, all tables have at least one partition, whether or not a partition was explicitly defined on the table. When you upgrade a database to Adaptive Server Enterprise 15.x, all existing tables in the database are modified to have one partition. When data is accessed from a table, necessary partitions are kept open; however, a DBA can limit the number of open partitions at any given point of time. This error indicates that there is not enough room to open a new partition.
Action	Use <code>sp_configure</code> to set the number of open partitions value appropriately for your server. See “Setting Configuration Parameters” in the <i>System Administration Guide: Volume 1</i> for details on setting this parameter.
Version in which this error is raised	Version 15.0 and later

Error 14313

Severity	11 or 21
Message text	Partition condition for object %ld, indid %d, ptncondid %ld was not found in database %d.

Explanation	<p>Semantic partitioned (range, hash, and list partitions) tables have a partition condition associated with the table's schema. This condition is stored as a partition condition object in <code>sysprocedures</code> and is linked to the source partitioned table. This partition condition object is used by the database task for access (reads and writes) of the table's partitions.</p> <p>When a database task fails to retrieve the partition condition object associated with a partitioned table, error 14313 is raised. This failure can be an internal error (severity 21) or the task might be trying to retrieve information about a non-existent object (severity 11).</p>
Action	<p>This error can result due to problems encountered when creating or dropping a partition or a partition (local) index. If a different error occurs when a command is executed, correct that condition and retry the command. If no other way is found to correct the problem, it may be necessary to unpartition and repartition the table.</p>
Additional information	<p>If the error persists, it may be necessary to gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section "Other Useful Tasks" in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.</p>
Version in which this error is raised	<p>Version 12.5.1 and later</p>

The `mount` command is used to attach a database to a destination or secondary Adaptive Server. The `mount` command decodes the information in the *manifest* file and makes the set of databases available online. This chapter describes errors that can occur when you use the `mount` command.

Error	Page
Error 14503	623
Error 14504	624
Error 14511	624
Error 14513	625
Error 14560	626
Error 14561	626

Error 14503

Severity	20
Message text	I/O error occurred while reading '%.*s' as a manifest file. The file is not a manifest file or it is corrupt.
Explanation	A <i>manifest</i> file is used to store information about the database and devices used when quiescing or unmounting and mounting databases. Error 14503 is raised when Adaptive Server is unable to read a record from the <i>manifest</i> file.
Action	<ul style="list-style-type: none"> • Verify that the specified file is the <i>manifest</i> file created during the unmount or quiesce database operation. • Verify that the Adaptive Server user has permission to read the <i>manifest</i> file. <p>A <i>manifest</i> file has a binary format, and if you use FTP transfer this file without specifying the binary option, it may become corrupted. FTP the <i>manifest</i> file again, specifying binary, and reissue the command for which error 14503 was been raised.</p>

Additional information	If error 14503 persists, it may be necessary to gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.5.1 and later

Error 14504

Severity	16
Message text	<code>I/O error occurred while writing the manifest file to `%. *s`.</code>
Explanation	A <i>manifest</i> file is used to store information about the database and devices used when quiescing and/or unmounting and mounting databases (see “Manifest file” in the chapter “Database Mount and Unmount,” of the most recent version of the <i>System Administration Guide: Volume 2</i>). Error 14503 is raised when Adaptive Server is unable to write a page to a <i>manifest</i> file.
Action	<ul style="list-style-type: none">• Verify that the Adaptive Server user has permission to write to the <i>manifest</i> file directory.• Verify that there is sufficient free space in that directory to hold the <i>manifest</i> file.• If those conditions are correct, check the operating system log for any indications of problems with I/O to the device in question.
Additional information	If error 14503 persists, gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.5.1 and later

Error 14511

Severity	17
Message text	<code>** %s: Failed to allocate %d bytes of memory.</code>

Explanation	For more information about mount and dismount, see “Database Mount and Unmount,” in the most recent version of <i>System Administration Guide: Volume 2</i> . For details about quiesce database (added in version 15.0.2), see “Creating Databases and Tables,” in the most recent version of the <i>Transact-SQL User’s Guide</i> .
Action	Increase the procedure cache size for the appropriate Adaptive Server, or wait and try the command again when there are fewer processes using available memory.
Version in which this error is raised	Version 12.5.3 and greater

Error 14513

Severity	16
Message text	Cannot parse manifest file data. Manifest file '%.*s' is corrupt.
Explanation	While parsing the <i>manifest</i> file, Adaptive Server validates each <i>manifest</i> item type it encounters. When the item type validation fails, error 14513 is raised.
Action	<ul style="list-style-type: none"> • Verify that the specified file is a <i>manifest</i> file. • If a <i>manifest</i> file is sent using FTP, without specifying binary mode, the file can become corrupted. Re-send the <i>manifest</i> file using FTP and specify binary mode.
Additional information	<p>If Adaptive Server still raises error 14513 while attempting to read the <i>manifest</i> file after resending it via FTP with the binary setting; or if the <i>manifest</i> file has never been sent using FTP, gather a configurable shared memory dump (CSMD) and contact Sybase Technical Support. Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.</p> <p>See</p>
Version in which this error is raised	Version 12.5.1 and later

Error 14560

Severity	20
Message text	Attempt to write the manifest file failed; error code is %ld.
Explanation	<p>A <i>manifest</i> file describes a database and is created by executing any of these commands:</p> <ul style="list-style-type: none">• quiesce database for external dump• unmount• create manifest file <p>While attempting to write <i>manifest</i> data from in-memory structure to a file, an error occurred.</p>
Action	<p>Check that the necessary file space exists and that the user has permissions to write to the file. Often there is another error that accompanies error 14560. If the failed command was <code>quiesce database for external dump</code>, it is possible that the database has already been quiesced even if the <i>manifest</i> file creation failed. In this case, issue the <code>quiesce database</code> command to resume activity on the database:</p> <pre>quiesce database <tag_name> release</pre> <p>Where <i><tag_name></i> is a user-defined name that designates the list of databases to hold or release. The <i><tag_name></i> must conform to the rules for identifiers.</p> <p>Resolve the other error first. Then retry <code>quiesce database for external dump</code> after the accompanying error is corrected.</p>
Version in which this error is raised	Version 12.5.1 and later

Error 14561

Severity	20
Message text	Attempt to read the manifest file failed; error code is %ld.
Explanation	<p>A <i>manifest</i> file is written to describe a database by the commands <code>quiesce database for external dump</code>, <code>unmount</code>, or <code>create manifest file</code>. While attempting to read the <i>manifest</i> file into memory, an error occurred.</p>

Action Possible values for the error code include:

Error code	Description	Explanation or action
1	Page has not been allocated.	Recreate the invalid <i>manifest</i> file.
2	Page header is invalid or corrupt.	Recreate the invalid <i>manifest</i> file.
3	Page data length is invalid.	Adaptive Server may have attempted to load a dump of a different page size.
4	Invalid service function ID.	Internal Adaptive Server Enterprise error. Contact Sybase Technical Support.
5	Could not allocate memory.	Insufficient memory available; increase procedure cache size and retry.
6	No method available for this task.	Recreate the invalid <i>manifest</i> file.
8	General read failure.	A possible corrupted <i>manifest</i> file or hardware error.
9	General failure to free memory.	Internal Adaptive Server Enterprise error. Contact Sybase Technical Support.
10	Page read is not a first header page.	Recreate the invalid <i>manifest</i> file.
11	Data ended before last header page read.	Recreate the invalid <i>manifest</i> file.
12	User method returned failure status.	Internal Adaptive Server Enterprise error. Contact Sybase Technical Support.
13	User queue ID is invalid.	Internal Adaptive Server Enterprise error. Contact Sybase Technical Support.

Retry the mount command after the error is corrected.

Additional information If the error code indicates an internal Adaptive Server Enterprise error, a shared memory dump may be necessary. Before contacting Sybase Technical Support, gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide.

Version in which this error is raised Version 12.5.1 and later

System catalogs use row anchors that are stored in a per-database pseudo-catalog named `sysanchors`. This chapter describes resolutions to errors that can occur when using row anchors.

Error	Page
Error 14605	629
Error 14618	630
Error 14619	630
Error 14620	631
Error 14621	631
Error 14622	632
Error 14623	632

Error 14605

Severity	21
Message text	The object id to which logical page 24 belongs to is set incorrectly as '%ld'. The expected object id is '%d'.
Explanation	Like <code>sysgams</code> , <code>sysanchors</code> is not a normal database table and cannot be queried. Adaptive Server attempted to write out an anchor page and at this time page 24 is expected to belong to <code>sysanchors</code> , but it does not.
Action	This is an internal error. Contact Sybase Technical Support.
Additional information	Before contacting Sybase Technical Support, gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 14618

Severity	21
Message text	There is not enough space on anchor page %d to insert a %d byte anchor. Space for this anchor should have been reserved, but was not.
Explanation	Adaptive Server attempted to redo or undo an anchor page change and attempted to insert an anchor. Because this action is a redo or undo, there should be space on this anchor page for the insert. but there is not.
Action	This is an internal error. Contact Sybase Technical Support.
Additional information	Before contacting Sybase Technical Support, gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 14619

Severity	21
Message text	The anchor for row ID '%S_RID' is missing from sysanchors page %d.
Explanation	<p>This is an internal error raised when Adaptive Server is recovering row anchor information on the anchor page. Row anchor information consists of the location of the system catalog row in the data page.</p> <p>The recovery process (start-time recovery or a task undergoing runtime recovery) reads the log records and applies changes to the page, which involves either redo or undo changes in the log record. In any case, during the recovery, row anchor information should persist on the anchor page involved in the log record, otherwise this is an internal error.</p>
Action	This is an internal error. Contact Sybase Technical Support.
Additional information	Before contacting Sybase Technical Support, gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 14620

Severity	21
Message text	An attempt to read page %d in database '%.*s' failed.
Explanation	This is an internal error raised when Adaptive Server attempts to redo or undo an anchor page and can not get the affected anchor page. The recovery process (boot-time recovery or a task undergoing runtime recovery) reads the log records and applies the changes to the page, which involves either a redo or undo of the changes in the log record. If the recovery process is not able to find the anchor page involved in the log record, it raises this internal error.
Action	This is an internal error. Contact Sybase Technical Support.
Additional information	Before contacting Sybase Technical Support, gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 14621

Severity	21
Message text	The anchor for database ID %d, start page %d is missing from sysanchors page %d.
Explanation	This is an internal error raised while Adaptive Server is recovering usage anchor information on the anchor page. Usage anchors cache the <code>sysusages</code> entries of the master and master companion database. The recovery process (start-time recovery or a task undergoing runtime recovery) reads the log records and applies the changes to the page, which involves either a redo or undo of the changes in the log record. If the recovery process is not able to find the anchor page involved in the log record, it raises this internal error.
Action	This is an internal error, please contact Sybase Technical Support.
Additional information	A configurable shared memory dump (CSMD) for error 14621 may help diagnose this problem. Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later.

Error 14622

Severity	20
Message text	<pre>The dbinfo anchor in the database '%S_DBINFO' is already locked by the task. Nested access to dbinfo anchor is not allowed. This is an internal error. Please contact Sybase Technical Support.</pre>
Explanation	This is an internal error raised when Adaptive Server is updating the “on disk database information” (dbinfo). The dbinfo structure consists of the database name, the status of the database, a pointer to the last checkpoint record, the first page of syslogs, and so on. If a database task tries to modify dbinfo information, the task must acquire a lock on the dbinfo anchor. This error is raised when a task tries to lock an already locked dbinfo anchor that was locked by the same task.
Action	This is an internal error. Contact Sybase Technical Support.
Additional information	Before contacting Sybase Technical Support, gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

Error 14623

Severity	20
Message text	<pre>You are trying to remove an anchor for a row in %s that still exists. You must delete the row before you may remove its anchor.</pre>
Explanation	This is an internal error that is raised when removing a row anchor. A row anchor gives the location of a system catalog row in the data page. If any task tries to remove a row anchor pointing to an active row, error 14623 is raised.
Action	This is an internal error. Contact Sybase Technical Support.
Additional information	Before contacting Sybase Technical Support, gather a configurable shared memory dump (CSMD). Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Troubleshooting and Disaster Recovery</i> guide.

Version in which this
error is raised

Version 15.0 and later

XML Utility Manager Errors (14900s)

Adaptive Server Enterprise version 12.5.0.1 and later provides the ability to handle semi-structured and unstructured XML data. This chapter describes errors that may occur when you work with XML documents. Adaptive Server Enterprise.

Refer to the *Adaptive Server Enterprise XML Services* guide for details.

Error	Page
Error 14901	635

Error 14901

Severity	20
Message text	Unexpected exception caught at %s:%d.
Explanation	The XML manager has encountered a problem at the indicated source file and line. This error cannot be handled by the normal error handling routines.
Action	This is an internal Adaptive Server error. Validate the XML documents being used and contact Sybase Technical Support.
Additional information	If the error persists, it may be necessary to gather a configurable shared memory dump (CSMD) before contacting Sybase Technical Support. Information about CSMD can be found in the section “Other Useful Tasks” in the <i>Adaptive Server Enterprise Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.5.1 and later

Adaptive Server Enterprise version 12.5.2 and later includes messaging functionality with the real time data services (RTDS) option package. This option simplifies the development of applications that interact with messaging systems and databases. This chapter describes errors that may occur when using this feature. See the *Adaptive Server Enterprise Messaging Services User's Guide* for more information.

Error	Page
Error 15101	637
Error 15105	638

Error 15101

Severity	17
Message text	There is insufficient memory to allocate %ld bytes. Please increase configuration parameter `messaging memory`, or try again when there is less activity on the server.
Explanation	The real-time data services (RTDS) manager attempted to allocate memory for a message but was unable to do so.
Action	If the number of bytes to be allocated is higher than expected, check the application to make sure the message sizes are not excessive. If the state is 1 or 3, increase “messaging memory.” If the error is raised with state 2, increase the “heap memory per user.”
Additional information	If the error persists after you make configuration changes, a configurable shared memory dump (CSMD) may be helpful in determining the root cause. Refer to “Obtaining a CSMD dump” in the “Other Useful Tasks” chapter of the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.5.2 and later

Error 15105

Severity	20
Message text	Messaging services API call failed. Operation '%1!'
Explanation	The connection using the Adaptive Server Enterprise real-time messaging service (RTMS) has failed.
Action	<p>Typically error 15106 follows this error, which provides additional information from the messaging provider. Address this error first, and check that the message being sent is syntactically correct.</p> <p>If the error persists, contact Sybase Technical Support.</p>
Version in which this error is raised	Version 12.5.2 and later

Web Services Errors (15200s)

Adaptive Server Enterprise Web Services uses the Web Services Engine that runs independently of Adaptive Server Enterprise. This chapter describes errors that may occur as a result of using Web Services.

Error	Page
Error 15219	639

Error 15219

Severity	20						
Message text	Cannot retrieve object ID for `%.*s`.						
Explanation	The Web Services Manager tried to access an object that should exist, but does not. Error 15219 is raised with the following states:						
	<table border="1"> <thead> <tr> <th>State</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>The proxy table entry for the specified Web service does not exist in the <code>sysobjects</code> table.</td> </tr> <tr> <td>2</td> <td>The “webservices_role” does not exist in the system.</td> </tr> </tbody> </table>	State	Meaning	1	The proxy table entry for the specified Web service does not exist in the <code>sysobjects</code> table.	2	The “webservices_role” does not exist in the system.
State	Meaning						
1	The proxy table entry for the specified Web service does not exist in the <code>sysobjects</code> table.						
2	The “webservices_role” does not exist in the system.						
Action	<ul style="list-style-type: none"> For state 1, verify that the specified Web service proxy table is present in <code>sysobjects</code>. If necessary, re-add the service using <code>sp_webservices</code>. For state 2, re-add the “webservices_role.” 						
Additional information	Refer to the <i>Adaptive Server Enterprise Web Services User’s Guide</i> for more information about implementing Web services. If the recommended actions do not correct the error, contact Sybase Technical Support.						
Version in which this error is raised	Version 15.0 and later						

Encryption Errors (15400s)

This chapter describes errors that occur as a result of using encryption in Adaptive Server Enterprise. Refer to the *System Administration Guide: Volume 1* for more information about using security and data encryption.

Error	Page
Error 15402	641
Error 15403	642
Error 15412	642
Error 15413	643
Error 15417	643

Error 15402

Severity	18
Message text	The encrypted columns feature is not available on this platform.
Explanation	The encrypted columns feature was introduced with Adaptive Server Enterprise 12.5.3a for many major platforms. More platforms have been added with new versions of Adaptive Server Enterprise. Error 15402 is raised if you are running on a platform and Adaptive Server Enterprise version combination that does not support encrypted columns.
Action	Encryption may not be available feature on your platform for a specific Adaptive Server Enterprise release. Check the version-specific release bulletin section “Product and platform interoperability.”
Version in which this error is raised	Version 12.5.3a and later

Error 15403

Severity	17
Message text	There is no memory available for an internal encryption operation in Adaptive Server. The current command cannot be processed.
Explanation	When Adaptive Server attempted an encryption operation, the Server also tried to allocate memory and was unable to do so.
Action	The encryption memory pool is a fixed size and cannot be increased. Retry the failing query when the server is less busy.
Additional information	If the error persists, a configurable shared memory dump (CSMD) may be helpful in determining the root cause. Information about CSMD can be found in the chapter “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.5.3a and later

Error 15412

Severity	18
Message text	An internal error occurred when Adaptive Server was setting up the context for symmetric encryption or decryption (error code %1!). Please contact Sybase Technical Support.
Explanation	While creating, retrieving, or applying an encryption key, Adaptive Server detected an error.
Action	Turning on trace flag 15401 to print additional information about the error. If that information is not sufficient to help diagnose the problem, contact Sybase Technical Support.
Additional information	A configurable shared memory dump (CSMD) may be helpful in determining the root cause. Information about CSMD can be found in the chapter “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 12.5.3a and later

Error 15413

Severity 18

Message text An internal error occurred during an encryption operation (error code %d). Please contact Sybase Technical Support.

Explanation Error 15413 is raised with these states:

State	Meaning
2	The encryption function failed.
3	The computed pad length of an encryption buffer was < 0.
4	While attempting to encrypt a column; the internal buffer length allocated was less than the size needed to encrypt.
5	While attempting to decrypt a column, the internal buffer length allocated was less than the size needed.
6	The destination for an encrypted column is too small to hold the encrypted value.
7	The destination for a decrypted column is too small to hold the decrypted value.

Action Error 15413 is an internal Adaptive Server error. Please have the description of the encrypted columns in use ready and contact Sybase Technical Support.

Additional information If the error persists, a configurable shared memory dump (CSMD) may be helpful in determining the root cause. Information about CSMD can be found in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide.

Version in which this error is raised Version 12.5.3a and later

Error 15417

Severity 18

Message text An internal error occurred during a decryption operation (error code %d). Please contact Sybase Technical Support.

Explanation Error 15417 is raised with these states:

State	Meaning
1	The decryption function failed.
2	The attempt to decrypt the concatenation of password and salt into a local buffer failed.
3	The attempt to encrypt key and salt (plus any login association data) into a buffer failed.
4	The attempt to decrypt a password failed.

Action Error 15417 is an internal Adaptive Server error. Please have the description of the encrypted columns in use ready and contact Sybase Technical Support.

Additional information If the error persists, a configurable shared memory dump (CSMD) may be helpful in determining the root cause. Information about CSMD can be found in the chapter “Other Useful Tasks” in the most recent version of the *Troubleshooting and Disaster Recovery* guide.

Version in which this error is raised Version 12.5.3a and later

This chapter describes errors that occur related to Adaptive Server Enterprise's gathering and managing of statistics. Refer to *Performance and Tuning: Monitoring and Analyzing* for more information about gathering statistics and the `systabstats` and `sysstatistics` tables, which store statistics for all tables, indexes, and any unindexed columns for which you have explicitly created statistics.

Error	Page
Error 15600	645
Error 15601	646

Error 15600

Severity	20
Message text	Lookup of systabstats for %S_PTINFO failed.
Explanation	<p>A row for a given object ID, index ID, and partition Id was not found in <code>systabstats</code>. This may happen during one of these scenarios:</p> <ul style="list-style-type: none"> • While attempting to obtain <code>systabstats</code> for a given object ID, index ID, and partition ID to place the statistics data in a partition descriptor, no row was found.; or • While attempting to delete a row in <code>systabstats</code> for a given object ID, index ID, and partition ID, the row to be deleted was not found.
Action	This is an Adaptive Server internal error. Please contact Sybase Technical Support.
Additional information	If the error persists, a configurable shared memory dump (CSMD) may be helpful in determining the root cause. Information about CSMD can be found in the chapter "Other Useful Tasks" in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15 and later

Error 15601

Severity	20
Message text	Failed to cache systabstats contents for %S_OBJINFO.
Explanation	When an object is opened, a descriptor is built and put into the descriptor cache. As part of that process, the statistics for the object are read from <code>systabstats</code> and cached in the object's partition descriptor. Error 15601 is raised when is an error caching the statistics.
Action	<p>Each object must have at least one valid row in <code>systabstats</code>. Error 15601 may be raised when no such row exists. If it appears the proper rows exist in <code>systabstats</code>, verify that the values are high enough for the configuration properties “number of open indexes,” “number of open objects,” and “number of open partitions.”</p> <p>If the 15601 is raised during an online database following a load, run update statistics on the object in the source database and re-dump it.</p>
Additional information	If the error persists, a configurable shared memory dump (CSMD) may be helpful in determining the root cause. Information about CSMD can be found in the chapter “Other Useful Tasks” in the most recent version of the <i>Troubleshooting and Disaster Recovery</i> guide.
Version in which this error is raised	Version 15.0 and later

System Procedure Errors (17000s, 18000s)

This section contains error messages raised by Adaptive Server system procedures.

Note When system procedure error messages are raised, the error number displays as 0 (zero) and the global variable @@error returns 0. This section includes non-zero error numbers along with the error message text. If you call Sybase Technical Support about an error in this section, include the error number that is shown here even though that number did not display when the error was raised.

Error	Page
Error 17461	648
Error 17715	649
Error 17716	649
Error 17737	650
Error 17870	650
Error 17871	651
Error 17872	651
Error 17873	652
Error 17874	652
Error 17875	653
Error 17903	653
Error 17904	654
Error 17905	654
Error 17906	654
Error 17910	655
Error 18031	655
Error 18032	656

Error	Page
Error 18033	657
Error 18123	657
Error 18750	658

Error 17461

Severity	0
Message text	Object does not exist in this database.
Explanation	<p>This error occurs when the object you specified does not exist in the current database when you run any of the following stored procedures:</p> <ul style="list-style-type: none">• sp_auditobject• sp_auditsproc• sp_depends• sp_estspace• sp_foreignkey• sp_help• sp_helppartition• sp_helpconstraint• sp_helpindex• sp_helptext• sp_primarykey• sp_procqmode• sp_remap• sp_spaceused
Action	Make sure you are in the database in which the object resides before entering the stored procedure command. Use the owner name in your object name specification if you are not the object owner.
Version in which this error is raised	All versions

Error 17715

Severity	0
Message text	The last-chance threshold for database %1! is now %2! pages.
Explanation	This informational message displays when a last-chance threshold has been established on a newly created log device (<code>sp_logdevice</code>), or an existing log device has been adjusted after a procedure was executed (<code>sp_dropsegment</code> , <code>sp_extendsegment</code>) that alters the size of a log device. The threshold is established automatically by Adaptive Server.
Action	No action is necessary.
Version in which this error is raised	All versions

Error 17716

Severity	0
Message text	Could not update the last-chance threshold for database %1!
Explanation	<code>lct_admin</code> runs automatically when any attempt is made to add or modify the log segment of a database with <code>sp_logdevice</code> , <code>sp_dropsegment</code> , or <code>sp_extendsegment</code> . Error 17716 occurs when an attempt to run such a command fails.
Action	Check other messages that display near the attempted command. They may give you additional information regarding the failure of <code>lct_admin</code> or may indicate the original command was not successful. If the command appeared to run successfully, except for this error message, contact Sybase Technical Support about evaluating the current status of the last-chance threshold in this database.
Additional information	Before calling Technical Support, have the information available that is listed in “Reporting errors” on page 4.
Version in which this error is raised	All versions

Error 17737

Severity	0
Message text	Partitioned objects cannot be moved.
Explanation	Adaptive Server does not allow you to use <code>sp_placeobject</code> on partitioned objects. Error 17737 occurs when you try to use <code>sp_placeobject</code> to change future space allocations on a partitioned table.
Action	To use <code>sp_placeobject</code> on the table: 1 Use the <code>unpartition</code> clause of the <code>alter table</code> command to concatenate all partitions: <pre>1> alter table <table_name> unpartition 2> go</pre> 2 Run the <code>sp_placeobject</code> command again.
Additional information	Refer to the <i>Reference Manual: Commands</i> for information about <code>alter table</code> and the <i>Reference Manual: Procedures</i> for details about <code>sp_placeobject</code> .
Version in which this error is raised	All versions

Error 17870

Severity	0
Message text	Table '%1!' does not exist in database '%2!' -- cannot add thresholds.
Explanation	<code>systhresholds</code> should exist in every database in Adaptive Server. Error 17870 occurs when the system table <code>systhresholds</code> does not exist in the current database, and <code>sp_addthreshold</code> or <code>sp_modifythreshold</code> was not able to add or modify the requested threshold.
Action	Run <code>dbcc checkalloc</code> and <code>dbcc checkdb</code> to determine if the database is corrupt. Contact Sybase Technical Support if <code>systhresholds</code> is missing from a database. When you call Technical Support, have the information ready that is listed in “Reporting errors” on page 4, including the output of: <ul style="list-style-type: none">• <code>select @@version</code>• <code>dbcc</code> commands

Additional information	Refer to “sp_addthreshold” in the <i>Reference Manual: Procedures</i> for additional information.
Version in which this error is raised	All versions

Error 17871

Severity	0
Message text	There is no segment named '%1!'.
Explanation	This error occurs when an attempt is made to add, drop, or modify (using sp_addthreshold, sp_droptreshold, or sp_modifythreshold) a segment that does not exist in the current database.
Action	Run sp_helpsegment to display segment information about the current database. If the segment exists, check that you have spelled it correctly in your procedure. If the segment does not exist and you believe it has been dropped, use sp_addsegment to add it back. If the requested segment does not exist and you are certain no one has dropped it, run dbcc commands to check the integrity of the database.
Additional information	Refer to “sp_addthreshold” in the <i>Reference Manual: Procedures</i> and “dbcc” in <i>Reference Manual: Commands</i> for additional information.
Version in which this error is raised	All versions

Error 17872

Severity	0
Message text	This threshold is too close to one or more existing thresholds. Thresholds must be no closer than 128 pages to each other.
Explanation	This error occurs when an attempt is made to add or modify thresholds (using sp_addthreshold or sp_modifythreshold), that are too close together.
Action	Thresholds for a database cannot be closer than:

$(\langle \text{size_of_threshold} \rangle + (2 * @@\text{thresh_hysteresis}))$

Ensure that the free space allocation allows enough pages. Rerun the command with the appropriate value for free space.

Additional information Refer to the chapter “Managing Free Space with Thresholds” in the *System Administration Guide: Volume 2* for information about threshold size and hysteresis value.

Version in which this error is raised All versions

Error 17873

Severity 0

Message text Adding threshold for segment '%1!' at '%2!' pages.

Explanation This message displays after the successful completion of a `sp_addthreshold` command or a `sp_modifythreshold` command that adds a threshold.

Action No action is required.

Version in which this error is raised All versions

Error 17874

Severity 0

Message text A threshold at %1! pages is logically impossible for segment '%2!'. Choose a value between %3! and %4! pages.

Explanation This error occurs when an attempt is made to either add or modify a threshold (using `sp_addthreshold` or `sp_modifythreshold`) for a number of pages that is greater than the segment named in the command.

Action Run `sp_helpdb <database_name>` to display the device fragments used by the specified database and the size of the device fragment. Then rerun the command with an appropriate page value for `<free_space>`.

Version in which this error is raised All versions

Error 17875

Severity	0
Message text	Only the DBO of database '%1!' or a user with System Administrator (SA) role may add, delete, or modify thresholds in that database.
Explanation	This error occurs when you try to run <code>sp_addthreshold</code> , <code>sp_droptreshold</code> , or <code>sp_modifythreshold</code> and you are not the owner of the current database and do not have the “sa” role.
Action	<p>If you are not the database owner, run <code>sp_displaylogin</code> to display roles currently assigned to you. (You must have the <code>sso_role</code> or <code>sa_role</code> to display this information about another user.)</p> <p>Ask your system administrator to grant you the <code>sa_role</code> or ask another user to execute the command.</p>
Version in which this error is raised	All versions

Error 17903

Severity	0
Message text	Table '%1!' does not exist in database '%2!' -- cannot drop thresholds.
Explanation	<code>systhresholds</code> must exist in every database in Adaptive Server. Error 17903 occurs when the system table <code>systhresholds</code> does not exist in the current database, and <code>sp_droptreshold</code> is unable to drop the requested threshold.
Action	<p>Run <code>dbcc checkalloc</code> and <code>dbcc checkdb</code> to determine if there is corruption in the database.</p> <p>Contact Sybase Technical Support if <code>systhresholds</code> is missing from a database.</p>
Additional information	<p>When you call Technical Support, have the information ready that is listed in “Reporting errors” on page 4, including the output of:</p> <ul style="list-style-type: none">• <code>select @@version</code>• All error messages• <code>dbcc</code> commands

Version in which this error is raised All versions

Error 17904

Severity 0

Message text Segment '%1!' does not have a threshold at '%2!' pages.

Explanation This error occurs when an attempt is made to either drop or modify a threshold (using `sp_droptreshold` or `sp_modifythreshold`) that does not exist for the free space page allocation given in the command.

Action Run `sp_helpthreshold` without any options to display all the thresholds and the free page allocations for the current database, then rerun the stored procedure with corrected parameters.

Version in which this error is raised All versions

Error 17905

Severity 0

Message text You may not drop the log's last-chance threshold.

Explanation The last-chance threshold is created and maintained by the Adaptive Server and cannot be dropped by a user. Error 17905 occurs when an attempt is made to drop the database last-chance threshold with `sp_droptreshold`.

Action None.

Version in which this error is raised All versions

Error 17906

Severity 0

Message text	Dropping threshold for segment '%1!' at '%2!' pages.
Explanation	This message is issued after successful completion of an <code>sp_droptreshold</code> command or an <code>sp_modifythreshold</code> command that drops a threshold.
Action	None.
Version in which this error is raised	All versions

Error 17910

Severity	0
Message text	Database '%1!' has no thresholds -- table '%2!' does not exist.
Explanation	The <code>systhresholdssystem</code> table should exist in every database in Adaptive Server. Error 17910 occurs when <code>systhresholds</code> does not exist in the current database and <code>sp_helpthreshold</code> is unable to output information from that table.
Action	Run <code>dbcc checkalloc</code> and <code>dbcc checkdb</code> to determine if there is database corruption. If there is, contact Sybase Technical Support.
Additional information	When you call Sybase Technical Support, have the information ready that is listed in “Reporting errors” on page 4, including the output of: <ul style="list-style-type: none"> • <code>select @@version</code> • All error messages • <code>dbcc</code> commands
Version in which this error is raised	All versions

Error 18031

Severity	0
Message text	This procedure can only affect thresholds in the current database. Say 'USE %1!', then run this procedure again.

Explanation	Thresholds for a specified database can only be created or modified while in the database. Error 18031 occurs when you try to run <code>sp_addthreshold</code> , <code>sp_droptreshold</code> , or <code>sp_modifythreshold</code> when you are not in the database for which you are running the command.
Action	<ol style="list-style-type: none">1 Determine the name of the current database:<pre>1> select db_name() 2> go</pre>2 Select the correct database:<pre>1> use <database_name> 2> go</pre><p>Where <code><database_name></code> is the name of the database for which you want to run the threshold command.</p>3 Rerun the <code>sp_addthreshold</code>, <code>sp_droptreshold</code>, or <code>sp_modifythreshold</code> command.
Version in which this error is raised	All versions

Error 18032

Severity	0
Message text	You may not alter the free space or segment name of the log's last-chance threshold.
Explanation	The only parameter you can change for the last-chance threshold is the name of the procedure that will execute when the threshold is crossed. Error 18032 occurs when you try to modify the last-chance threshold name or free space allocation with <code>sp_modifythreshold</code> .
Action	None.
Version in which this error is raised	All versions

Error 18033

Severity	0
Message text	You cannot drop user because user '%1!' owns thresholds in database.
Explanation	The user who created a threshold or last modified a threshold is the threshold owner. A user who owns a threshold cannot be dropped. Error 18033 occurs when you use <code>sp_dropuser</code> to try to drop a user who owns a threshold.
Action	Use any of these three options: <ul style="list-style-type: none">• Drop the threshold procedure owned by the user, then drop the user; or:• Have another user with appropriate privileges modify the threshold, then drop the user; or:• To inactivate the user Adaptive Server login while leaving the threshold active, use <code>sp_locklogin</code> to prevent use of the login.
Additional information	See “ <code>sp_locklogin</code> ” in <i>Reference Manual: Procedures</i> for more information about using the <code>sp_locklogin</code> command.
Version in which this error is raised	All versions

Error 18123

Severity	0
Message text	Configuration option changed. The SQL Server must be rebooted before the change in effect since the option is static.
Explanation	When you use <code>sp_configure</code> to modify a static parameter: <ul style="list-style-type: none">• The configuration value is updated• The configuration file is updated• The change takes effect when you restart Adaptive Server

	Error 18123 is an informational message relating that a static parameter has been changed and that the change will not take effect until Adaptive Server is restarted.
Action	Shut down and restart Adaptive Server to make the parameter change take effect.
Additional information	Refer to “sp_configure” the <i>Reference Manual: Procedures</i> for information about using the sp_configure command.
Version in which this error is raised	All versions

Error 18750

Severity	0
Message text	The command '%1!' can not be run in this server. The server is currently configured with mode '%2!'.
Explanation	<p>In the Adaptive Server Enterprises High Availability (HA) cluster configuration, two nodes in a cluster can take over each other’s workload in the event of a failure; once the problem with the primary is resolved, you can use failback to move the primary companion’s shared disks from the secondary node back to the primary node and start the primary companion on the primary node.</p> <p>Error 18750 is raised if you try to assign a (primary or secondary) server to a mode that is incompatible with its current state. For example, the error is raised if you issue sp_companion 'prepare_failback' on the secondary companion, but the secondary is not in the correct state for failback.</p>
Action	<p>Take the following steps:</p> <ol style="list-style-type: none">1 Check the companion mode of your servers by issuing the following on both primary and secondary servers:<pre>1> select @@cmpstate 2> go</pre><p>This command returns a value denoting the current companion mode. Refer to “Determining the companion’s mode” in the chapter “Modes of Failover” in <i>Using Sybase Failover in a High-Availability System</i> for a description of the return values.</p>

- 2 Check that the command that raised the 18750 message (appearing in the message text) is compatible with the server's companion mode. For example, you cannot issue a `suspend` command to a server currently in failover mode.

There are situations when the 18750 message requires corrective action. For example, if you issued `prepare_failback` and your primary companion is in normal companion mode but the secondary companion is in secondary failover mode, you must recover from this inconsistent state manually. The process involves shutting down both servers, restarting the secondary and repairing any suspect databases on the secondary. The steps for achieving this are platform-specific.

Additional information

Refer to the “Troubleshooting” section in the configuration chapter that matches your platform in *Using Sybase Failover in a High-Availability System*.

Version in which this error is raised

Version 12.0 and later

Kernel Errors

This section includes error messages generated by the Adaptive Server kernel, which have text descriptions but no error numbers. These errors display only in the Adaptive Server error log.

These error messages are presented in alphabetical order, do not have a “severity” level, and often do not state the versions of Adaptive Server that can be affected.

Error	Page
Asynchronous I/O not available	662
Buffer mismatch	663
Could not create shared memory	664
Current process infected	665
dopen error	667
dstart I/O request repeatedly delayed	669
Failed to open device	672
Feature is not licensed	673
File already in use	674
Memory too fragmented	675
Memory usage in procedure headers	676
Network information message	677
No more alarms available	678
Open objects parameter may be too low	679
os_attach_region errors	680
os_create_region errors	681
Read/write errors	683
Server is unnamed	685
Stack guardword corrupted	686
time slice error	689
t_rcvconnect error	691
udunmirror errors	691
No more engines can be offlined	693

Asynchronous I/O not available

Message text	<code>basis_daiioconfig: Asynchronous I/O not available</code>
Explanation	<p>This error (or additional errors starting with <code>basis_daiioconfig</code>) are raised when Adaptive Server is unable to start up with asynchronous I/O on an HP-UX system. This problem is caused by an incorrect operating system configuration. Common reasons include:</p> <ul style="list-style-type: none">• The user starting the server belongs to a group that has inadequate shared memory locking privileges.• There is insufficient lockable shared memory.• System memory and swap space have become fragmented.• Asynchronous I/O is not available. <p>This section explains how to investigate and correct the error, and provides additional resources for troubleshooting the memory problem.</p>
Action	<p>Although they are reported in the Adaptive Server error log, <code>basis_daiioconfig</code> errors originate from the operating system and can be corrected with appropriate management of memory parameters and device permissions.</p> <p>Investigate the problem by checking for additional asynchronous I/O related messages in the error log; these messages are printed just prior to the <code>Asynchronous I/O not available</code> error.</p> <ul style="list-style-type: none">• An error like the following means that the operating system does not have enough shared lockable memory to allocate to asynchronous I/O buffers.<pre>ioctl(ASYNC_CONFIG) failed: Not enough space</pre>Solutions vary, depending on the available resources and the type of tuning you choose. Options include: <ul style="list-style-type: none">• Increase swap space. Ensure that the primary swap space is at least 1.5 times the amount of RAM.• Add more physical memory.• Restart the machine. If the space problem is due to memory fragmentation, a restart may solve the immediate problem (which will eventually reappear).• Reduce the server's memory usage. If you have multiple Adaptive Servers running on the machine, consider the use of all servers.

- Set the kernel parameter `SWAPMEM_ON` to 0 (zero), increasing the amount of shared lockable memory available.
- Examine the kernel parameter `UNLOCKABLE_MEM`. A value of 0 allows the operating system to dedicate 75 percent of physical memory to lockable memory. Consult the HP-UX documentation at <http://www.hp.com> for details.
- An error like the following indicates a permission problem.

```
ioctl(ASYNC_CONFIG) failed: Not owner
```

Check that the file named `/dev/async` exists and is owned by user “sybase”, using this command:

```
ls -al /dev/async
```

It is also possible that the group to which the user starting the server belongs may not have permission to lock shared memory (`MLOCK` privilege). For details and instructions on setting the correct permissions, refer to the Sybase TechNote 20380 in the Sybase Technical Library at <http://www.sybase.com/detail?id=20380>.

Additional information

Refer to the following information about asynchronous I/O:

- *Adaptive Server Enterprise Installation Guide for HP-UX*. The “Pre-installation tasks” section contains information about enabling asynchronous I/O.
- *Permissions Information for Upgrading to Adaptive Server Enterprise 11.5 or 11.9 for HP-UX Customers*, in TechNote 20380 at <http://www.sybase.com/detail?id=20380>.
- *Release Bulletin Replication Server for HP-UX*. The “Special installation instructions” section describes how to enable asynchronous I/O for Replication Server.

Buffer mismatch

Message text

```
Bufunhash: Buffer %lx from database '%.*s' has page
number %ld in the page header and page number %ld in
the buffer header. Buffer will be unhashed, no action
```

Could not create shared memory

is necessary. Printing headers for diagnostics:

Note This error may be caused by a hardware problem.

Explanation This error only appears in the error log and means that a cache management problem occurred.

Warning! This can be a very serious error that is often followed by database corruption, such as 605 errors.

Although often a result of hardware failure, this error can also be caused by operating system or Adaptive Server problems.

Action

- 1 If possible, shut down and restart Adaptive Server immediately after this error occurs. This may prevent the buffer cache error from being flushed to disk.
- 2 Run complete diagnostics on the machine running Adaptive Server and on all disk drives and controllers attached to that machine.
- 3 Run complete `dbcc` checks, including `dbcc checkalloc` and `dbcc checkdb`, on any databases involved.
- 4 Repair or replace faulty hardware.
- 5 Shut down and restart Adaptive Server.
- 6 If you do not find any hardware problems, call Sybase Technical Support immediately.

Additional information Before calling Sybase Technical Support for assistance, have the information available that is listed in “Reporting errors” on page 4, including a reproducible case (if possible).

Could not create shared memory

Message text `kistartup: could not create shared memory`

Note This error may be caused by a hardware problem.

Explanation	<p>kistartup is the Adaptive Server kernel initialization routine that performs these tasks:</p> <ul style="list-style-type: none"> • Reads the configuration parameters • Allocates kernel tables and buffer pools • Sizes memory • Initializes devices • Spawns the service processes • Starts the scheduler <p>When kistartup is unable to create shared memory for kernel initialization, this error is written to the Adaptive Server error log. kistartup errors occur after <code>os_create_region</code> errors.</p>
Action	<p>Refer to “<code>os_create_region</code> errors” on page 681 for information about recovering from the <code>os_create_region</code> errors that precede the kistartup error. Once those errors are corrected, the kistartup error will not occur.</p> <p>Interaction with Monitor Server</p> <p>Shared memory errors can occur if you shut down Adaptive Server and restart it while Monitor Server is still active. To correct this type of error, shut down Monitor Server before you restart Adaptive Server.</p> <p>The interaction between Adaptive Server and Monitor Server is covered fully in the <i>Monitor Server User's Guide</i>.</p>

Current process infected

Message text `current process (0x%x) infected with %d`

Note This error may be caused by a hardware problem.

Explanation Adaptive Server reports this error when it detects a UNIX signal that specifies an error. The values (“%d”) that display in this error message vary by platform and Adaptive Server Enterprise versions; the most common values are 10 and 11.

Current process infected with 10

A value of 10 [SIGBUS] means that the operating system detected an address alignment error or a miscellaneous hardware error (for example, bus timeout).

A timeout can occur when the CPU issues a request across the bus for the contents of a memory location, and that request is not answered within that CPU's timeout period (usually a few nanoseconds).

Current process infected with 11

A value of 11 [SIGSEGV] means that the operating system detected a segment violation error.

Note Sometimes this error occurs in conjunction with a stack overflow or data corruption. For more information on stack overflow, refer to the error write-up "Stack guardword corrupted" on page 686.

The error message appears in the Adaptive Server error log followed by a stack trace. The "SQL causing error" or the *<lasterror>* that displays in the Adaptive Server error log may be the underlying cause for this error. But the message can also be just the last data Adaptive Server had in its cache space.

To identify the *<lasterror>* (except in the cases where the *<lasterror>* is 0), get the number that Adaptive Server displays by in the *<lasterror>* field from the Adaptive Server error log and consult this manual for more information on the error number.

In the following example, the value for *<lasterror>* is 614.

```
00: 94/02/14 11:32:26.02 kernel: current process (0x1fb001d)
infected with 11
00: 94/02/14 11:32:26.07 kernel: Address 0x808a6ef
(closetable+0x2f7), siginfo (code, address) = (2, 0x30)
00: 94/02/14 11:32:26.07 kernel: *****
00: 94/02/14 11:32:26.07 kernel: "SQL causing error" : CREATE TRIGGER
00: 94/02/14 11:32:26.07 kernel: curdb = 22 pstat = 0x10018
"lasterror = 614"
```

- Action
- 1 Try to eliminate the *<lasterror>*, which may be one of the causes for this error (except when *<lasterror>* is 0).
 - 2 Rerun the command referenced by the SQL causing error to see if the problem reoccurs.

If the process is infected with 11 and you can reproduce the problem, correct it as follows:

- If the SQL causing error is a compiled object such as a stored procedure, trigger, or view, drop and recreate the object.
- If the SQL causing error is ad hoc rather than a compiled object, moving the data may fix the problem. Use one of these options:
 - Select the table data into a new table, drop the old table, and rename the new table to the old table name;
 - or
 - Bulk copy the affected table out, drop and re-create the table, and bulk copy back in. This is the most efficient solution for a large table.

Note If moving the data corrects the problem, the data may have been corrupt. Be aware that moving corrupted data can lead to a data loss.

Check your hardware error log as this error can be caused by hardware failure as well.

Additional information

If this error persists, contact Sybase Technical Support. It would be useful to have a reproduction case (for example, a query that always causes the “infected” message) before calling Technical Support.

Before calling Technical Support for assistance, have the information hand that is listed in “Reporting errors” on page 4.

dopen error

Message text

```
dopen: open '%s', %s\n
```

Explanation

A dopen error is raised when Adaptive Server is unable to complete an asynchronous or synchronous disk I/O operation on a UNIX platform. For example:

```
00:1999/12/26 22:59:10.97 kernel dopen: open '/dev/vx/rdisk/db01076_m',
No such file or directory
```

dopen errors can occur during the Adaptive Server startup sequence, or during an engine startup. Depending on the nature of the problem, this error may be accompanied by other messages such as dopen errors or numbered errors. A dopen error aborts the startup sequence.

Action

Examine the server error log and note the exact `dopen` message text, the text immediately preceding and following the message, and any subsequent errors in the log. Corrective action depends on the type of problem encountered:

dopen: open '%s', no such file or directory

Make sure that the device, identified by the path shown in the message, exists and is correctly specified in the `RUN_<server>` file.

Note Each time a new Adaptive Server, Backup Server, or Monitor Server is created, the `srvbuild` program creates a `RUN_<server>` file that contains the information required to restart that server. `RUN_<server>` files are created in the `$$SYBASE/$SYBASE_ASE/install` directory. For Adaptive Server, the `RUN_<server>` file is named `RUN_<server_name>`, where `<server_name>` is the actual name of the server that was specified during installation.

For example, the `RUN_<server_name>` file name for an Adaptive Server named TEST is `RUN_TEST`.

For Backup Server and Monitor Server, the `RUN_<server_name>` files are named `RUN_<servername>_back` and `RUN_<servername>_mon`, respectively, where `<servername>` is the actual name of the server.

If the device no longer exists, restore it from known, clean backups.

dopen: open '%s', permission denied

Check permissions on the directory or disk file shown in the message. Verify that the account starting the server has the appropriate permissions.

dopen: open '%s', I/O error

Since Adaptive Server reports this error when its I/O operation does not receive a successful return value, this type of `dopen` error can be an early warning of disk problems. Examine the availability and condition of the device using your standard operating system procedures.

If you use mirroring, a `kdconfig` error following a `dopen` I/O error may indicate a mirror problem. For example:

```
00:99/05/26 16:05:29.64 kernel kdconfig: opening secondary master device
00:99/05/26 16:05:29.65 kernel dopen: open "", No such file or directory
00:99/05/26 16:05:29.65 kernel kdconfig: unable to read secondary master
device
```

Refer to “`udunmirror` errors” on page 691 for information about resolving mirror problems.

Additional information	When you call Sybase Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the <i>RUN_<server></i> file.
Version in which this error is raised	All versions

dstart I/O request repeatedly delayed

Message text kernel: dstartio: I/O request repeatedly delayed; block number: %ld, vdn: %ld

Note This error may be caused by a hardware problem.

Explanation This error message is not dangerous and there is no data corruption connected with it unless it is reported in the error log together with other kernel messages that indicate corruption.

Error message

If this is an error message, it is a sign of inconsistencies or corruption relating to the asynchronous I/O capabilities of Adaptive Server Enterprise. In this case, the above error message is reported along with other error messages in the Adaptive Server error log.

If this message is accompanied by other error or kernel messages in the Adaptive Server error log, refer to the appropriate section of this manual (by error number or name) to correct the additional errors.

Warning message (informational)

This can be an informational message (occurring by itself, any number of times). If this is a warning message, it is displayed by itself in the Adaptive Server error log. In this case, the message displays in the Adaptive Server error log when:

- Adaptive Server is configured for a different rate of concurrent asynchronous I/O requests than the operating system. Every time an I/O request has to be re-queued in Adaptive Server’s I/O queues, there is a decrease in overall performance due to re-queueing overhead.
- Adaptive Server is configured for a lower throughput rate of concurrent asynchronous I/O than it can handle.

Asynchronous disk I/O allows Adaptive Server to initiate a disk operation to continue processing other data while the disk operation is completed. This can improve performance and recoverability especially when the master device, the user databases, and transaction logs are on raw disk partitions. For more information on raw disk partitions, refer to the section “File system or raw partition” in the Adaptive Server Enterprise *Migration Guide* chapter “Making Database Administration Changes.”

Action

The following I/O-related Adaptive Server configuration parameters and operating system parameters play a role when this error occurs:

- `disk i/o structures` – the initial number of block I/O structures at Adaptive Server startup. This is an Adaptive Server configuration parameter.
- `max async i/os per server` – the number of asynchronous I/O requests outstanding at one time for Adaptive Server. This Adaptive Server configuration parameter applies only to multiprocessor machines.
- `max async i/os per engine` – the number of asynchronous I/O requests outstanding at one time for one engine (in the case of multiprocessor machines) and the maximum number of asynchronous I/O requests outstanding at one time for Adaptive Server (in the case of single processor machines). This is an Adaptive Server configuration parameter.
- Asynchronous I/O operating system parameters.

Determine the current I/O-related Adaptive Server configuration parameter settings by running the following commands:

```
1> sp_configure "disk i/o structures"
2> go
1> sp_configure "max async i/os per engine"
2> go
```

```
1> sp_configure "max async i/os per server"
2> go
```

Use the instructions in the sections below to determine appropriate values and then use `sp_configure` to change the values of `disk i/o structures`, `max async i/os per server`, and `max async i/os per engine`. In general:

- Set `max async i/os per server` to a value less than or equal to the total allowable I/O count for the system on which the server is running.
- Set `max async i/os per engine` to a value less than or equal to `max async i/os per server`.

AIX

Configure the operating system for a higher number of outstanding asynchronous I/Os using `smit`. Bring up the `smit aio` screen by typing `smit aio` at the AIX prompt. Then, go to the “Change/Show Characteristics of Asynchronous I/O” screen and raise the “maximum number of REQUESTS” parameter to 4096.

This change will improve Adaptive Server performance. You will not need to restart the operating system for the change to take effect. For more information on asynchronous I/O, search for “asynchronous” using the `info` utility.

Also, raise the values of the Adaptive Server I/O-related configuration parameters to 4096 to match the “maximum number of REQUESTS” value in `smit`.

Digital UNIX

In Digital UNIX releases before version 3.0 (Digital OSF), the number of outstanding asynchronous I/Os was hard-coded in the operating system kernel in the `AIO_MAX` parameter (stored in a header file in `/sys`). Its default value was 64. For those releases, the Adaptive Server I/O-related configuration parameters should be set to 64 so that Adaptive Server can match the operating system's rate of concurrent I/O requests.

In version 3.0 and later releases, you can change the number of outstanding kernel asynchronous I/Os. Refer to the Adaptive Server Enterprise installation and configuration guides for details.

Also, raise the values of the Adaptive Server I/O-related configuration parameters to match the value of the kernel's maximum outstanding asynchronous I/O requests.

HP-UX

Depending on the operating system kernel, the maximum number of outstanding asynchronous I/O requests can be either 50 (on older HP-UX 8.x kernels) or 500 (on patched HP-UX 8.x kernels and default on HP-UX on 9.x kernels). This value is stored in the `MAX_CONCURRENT` parameter in `/usr/include/sio/asynch`. Adjust the Adaptive Server I/O-related configuration parameters to match `MAX_CONCURRENT`.

The error log example below shows the Adaptive Server start-up messages associated with the older 8.x asynchronous I/O driver and its limit of 50 asynchronous requests at one time. Once a later asynchronous patch has been applied or the new operating system version has been installed, this limit is increased to 500.

```
00: 92/06/04 14:38:07.10 server: Number of blocks left
for proc headers: 351.
00: 92/06/04 14:38:07.36 kernel: basis_dioinit: Using older version
of asynch I/O driver with "50" concurrent requests
00: 92/06/04 14:38:07.38 server: Opening Master database...
```

Other operating systems

For other operating systems, check your operating system manuals for information about asynchronous configuration issues.

If problems persist, refer to the error documentation in this manual or contact Sybase Technical Support.

Additional information

Refer to:

- The section "File system or raw partition" in the Adaptive Server Enterprise Migration Guide chapter "Making Database Administration Changes"
- Your operating system documentation on asynchronous I/O

If you call Sybase Technical Support, have the information on hand that is listed in "Reporting errors" on page 4, including the output from `sp_configure` for I/O-related parameters or listing of your configuration file.

Failed to open device

Message text

```
kdopen: failed to open device %s for vdn %d
```

Note This error may be caused by a hardware problem.

Explanation

Before using a database device, Adaptive Server attempts to perform some initialization tasks. When mirroring is being used, if Adaptive Server is unable to open the secondary side of a mirrored device for these initialization tasks, this message is written to the error log. `<%s>` is the physical path to the secondary side of the mirrored device and `<%d>` is the virtual disk number of the device.

This error can be raised because:

- A database device is offline
- A database device has been removed or renamed

- A database file or device is inaccessible to Adaptive Server for any other reason such as insufficient permissions
- Action
- 1 Examine the Adaptive Server error log for other messages that may explain why this error has occurred. Refer to the instructions in this manual for resolving other errors associated with this error.
 - 2 Identify the device involved in the error:


```
1> select name from master..sysdevices
2> where low/power(2,24) = <vdevno>
```

Where <vdevno> is the virtual disk number of the device as shown in the error message.
 - 3 Examine the availability and condition of the device involved using your standard operating system procedures and make sure the device is accessible.

Feature is not licensed

- Message text `Feature %s is not licensed.`
- Explanation Sybase Software Asset Management (SySAM), which is available in Adaptive Server Enterprise version 12.0 and later, is used to manage licensed Sybase products and features. These features are licensed for customer use with paper-based SySAM certificates. Some examples of SySAM features include:
- ASE_SERVER permission to run the basic server itself
 - ASE_HA High Availability system features
 - ASE_JAVA Java in the database
- If a feature is not enabled, the message “Feature %s is not licensed” appears in the error log at the beginning of server startup, indicating no license was available. For example:
- ```
Feature ASE_JAVA is not licensed.
```
- Action                            Refer to the Sybase white paper *License Manager - Sybase Software Asset Management (SySAM)* available at <http://www.sybase.com/detail?id=1009350> for detailed instructions on resolving license issues.
- Additional information            • Adaptive Server error log
- License manager log – `$SYBASE/SYSAM-1_0/log/lmgrd.log`

- License file – `$SYBASE/SYSAM-1_0/licenses/license.dat`
- License certificates

Use the `sp_helpconfig` system stored procedure to see if a server option is enabled. For example, type the following procedure to display all configuration parameters containing “java” along with the parameter’s configured and runtime value.

```
sp_helpconfig java
```

Version in which this error is raised

Version 12.0 and later

## File already in use

Message text

```
basis_dlock: file '%s' already in use by a SQL Server
```

---

**Note** This error may be caused by a hardware problem.

---

Explanation

During startup, Adaptive Server tries to lock the master device. If Adaptive Server finds that the master device is already locked, the above message is written to the error log and Adaptive Server does not start. `<%S>` is the name of the master device file.

The most common reasons for this error are:

- Trying to start a server with a master device that is already being used by another Adaptive Server
- The server you are trying to start is already running

Action

At the operating system level, check the processes that currently exist and determine whether an Adaptive Server using the master device you specified is already running.

If an Adaptive Server using the specified master device is already running, and this is not the Server you want to start, determine the correct master device name for the Server you want to start.

If you cannot determine whether the master device is already in use, reboot the operating system.

---

**Note** If the master device resides on an operating system file system, copy the file to another name, modify the `RUN_server` file, and start Adaptive Server with the new file. See “Using a `RUN_server_name` file” in chapter “Starting and Stopping Servers” in the *Configuration Guide for UNIX*.

---

## Memory too fragmented

Message text `Physical memory on this machine may be too fragmented`

---

**Note** This error may be caused by a hardware problem.

---

Explanation At startup, Adaptive Server allocates and initializes the memory needed for its structures and cache. This error occurs when Adaptive Server is not able to initialize the memory needed for its structures and static cache.

The most common cause of this error is that the total memory required for the configurable parameters is greater than the memory configured for Adaptive Server.

Action **1** At the operating system level, check memory use.

On UNIX machines, there are several commands for checking memory use. Following are examples of two of these commands:

```
% /etc/pstat -s

44076k allocated + 8920k reserved =52996k used, 40960k available

% vmstat

procs memory page disk faults cpu
 r b w avm fre re at pi po fr de sr s0 d1 d2 d3 in sy cs us sy id
 4 0 0 0 336 0 1 0 0 0 16 0 1 0 0 0 181 260 399 84 2 14
```

If processes other than Adaptive Server are leaving memory fragmented, rebooting the operating system may resolve this error.

**2** Start Adaptive Server.

If you still cannot start Adaptive Server, refer to “Adaptive Server does not start after altering configuration” in the chapter “Adaptive Server System Recovery” in the most recent version of *Troubleshooting and Disaster Recovery*.

- 3 Once Adaptive Server is running, use `sp_configure` either to reduce the value of the configurable parameters (such as number of user connections and number of locks) or to increase the value of the total memory configuration parameter. The total memory parameter controls the size of memory, in 2K units, that Adaptive Server allocates from the operating system.

Refer to the Adaptive Server Enterprise performance and tuning guides and “Configuring Memory” in the *System Administration Guide: Volume 2* for information about memory-related configuration parameters.

Additional information

When you call Sybase Technical Support, have the information on hand that is listed in “Reporting errors” on page 4, including the text of your configuration file or `sp_configure` output.

## Memory usage in procedure headers

Message text

```
memory usage in procedure headers does not match memory
usage. count in pss (%d) for server process id %d.
```

Explanation

A process is a task being carried out by Adaptive Server. Processes can be in various states such as running or waiting for resources. For each user process and some system processes, Adaptive Server maintains a process status structure (PSS) to track the state of the process and the process’ use of resources. It also maintains an array of structures in the procedure header, each of which represents a physical page of memory in procedure cache.

When a process completes, Adaptive Server performs cleanup tasks such as freeing up resources used by the task. This error is reported when Adaptive Server cleans up after normal completion of a batch process, and determines that the memory usage tracked by the procedure header does not match the usage total maintained in the PSS.

This error may be due to an Adaptive Server problem.



|                                       |                                                                                                                                                                                                                                                                 |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Action                                | No action is required, since this is an informational message and does not result in any incorrect behavior. Upon detecting the memory count mismatch, Adaptive Server corrects the appropriate counter, completes the cleanup tasks, and continues processing. |
| Additional information                | If this error persists, contact Sybase Technical Support and have the information on hand that is listed in “Reporting errors” on page 4.                                                                                                                       |
| Version in which this error is raised | Versions 11.5, 11.0.3.1 and earlier                                                                                                                                                                                                                             |

## Network information message

|              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Message text | Using '%s' for network information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Explanation  | <p>Adaptive Server uses information from the network addressing (<i>interfaces</i>) file located in the <i>SYBASE</i> directory when it starts up and when it tries to establish a connection to a remote server. At those times, Adaptive Server displays this message, including the complete path of the interfaces file which was used.</p> <p>Following is an example of this message for UNIX:</p> <pre>kernel: Using '/usr/u/sybase/interfaces' for network information.</pre> <p>If a remote procedure call (RPC) is executed and no connection is currently established for that remote Server, Adaptive Server must first establish a connection to it, using a site handler process, before the RPC can be executed. Each time Adaptive Server establishes a site handler process, this informational message is written to the Adaptive Server error log.</p> <p>If the site handler process times out, this message can occur many times for each remote Server. If the site handler for a remote Server does not time out, then Adaptive Server writes this message for the first RPC to that remote Server, and the message will occur only once for each remote Server.</p> <p>Use <code>sp_helpserver</code> to check the status of the remote Adaptive Server. If the status column of the <code>sp_helpserver</code> output is “timeouts”, use the following action to keep the network information message to one occurrence.</p> |
| Action       | <p>To limit this message to one occurrence for a remote Server, disable the “timeouts” option to that remote Adaptive Server. Log in to <code>isql</code> as “sa” and use <code>sp_serveroption</code> to set the “timeouts” option to “no timeouts”:</p> <pre>1&gt; sp_serveroption &lt;remote_server&gt;, 2&gt; timeouts, false</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

```
3> go
```

Where *<remote\_server>* is the actual name of the remote server.

This action will keep the connection to the remote Server established and maintained until either the local or remote Server is shut down. If you have a limited number of available user connections, remember that one connection per remote Adaptive Server connection will be in constant use.

Additional information

Refer to the *Reference Manual: Commands* for information about using the `sp_serveroption` command.

## No more alarms available

Message text

```
uasetalarm: no more alarms available
```

Explanation

The Transact-SQL command `waitfor` defines a specific time, time interval, or an event for the execution of a statement block, stored procedure, or transaction. Adaptive Server uses alarms to carry out `waitfor` commands correctly.

The number of alarms needed by Adaptive Server is related to the requirements of an application and the number of instances of that application being run simultaneously. The maximum number of alarms available for use by Adaptive Server is controlled by the `number of alarms` configuration parameter.

When Adaptive Server needs more than the number of alarms available, this message is written to the error log.

Action

- 1 Determine the current value of the configuration parameter `number of alarms`:

```
1> sp_configure "number of alarms"
2> go
```

- 2 Change the value of `number of alarms` to a value larger than the current value:

```
1> sp_configure "number of alarms", <new_value>
2> go
```

## 3 Shut down and restart Adaptive Server.

---

**Note** Each alarm structure uses 20 bytes of memory. If you raise the value of number of alarms significantly, re-adjust Adaptive Server's memory use accordingly.

---

## Open objects parameter may be too low

|              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Message text | <pre>server: Warning: OPEN OBJECTS parameter may be too low; attempt was made to free up descriptor in desfree(). Run sp_configure to increase parameter value.</pre>                                                                                                                                                                                                                                                                                                                                                      |
|              | or                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|              | <pre>server: Warning: OPEN OBJECTS or OPEN DATABASES parameter may be too low; attempt was made to free up descriptors in release_dbtable(). Run sp_configure to increase parameter value.</pre>                                                                                                                                                                                                                                                                                                                           |
| Explanation  | <p>This message is a warning indicating that you do not have enough object descriptors or open databases available in the chain containing all the free descriptors.</p> <p>This message is displayed in the Adaptive Server error log when Adaptive Server reaches the end of its array of allocated object descriptors, tries to recycle previously-used object descriptors and fails. It can also be displayed when Adaptive Server has fewer open databases configured than the number of databases on the Server.</p> |
| Action       | <p>To recover from this error, increase the values of the configuration parameters number of open objects and number of open databases.</p> <p>The configuration parameter number of open databases sets the maximum number of databases that can exist at one time on Adaptive Server.</p> <p>The configuration parameter number of open objects sets the maximum number of objects that can be open at one time on Adaptive Server.</p>                                                                                  |

Increasing the value for number of open databases or number of open objects does not have a significant impact on performance or storage requirements, but it does require memory. For more information about how Adaptive Server allocates memory, refer to “Memory Use and Performance” in *Performance and Tuning: Basics* and “Configuring Memory” in the *System Administration Guide: Volume 2*.

Use the system procedure `sp_configure` to display the values of the configuration parameters you need to change:

```
1> sp_configure "number of open databases"
2> go
```

If the value for number of open databases is less than the value displayed by `select count (*) from sysdatabases`, increase the value of number of open databases. Otherwise, increase the value of number of open objects only.

To increase either of these values, refer to the instructions in Chapter 5, “Setting Configuration Parameters,” in the *System Administration Guide: Volume 1*.

Additional information

In Adaptive Server Enterprise version 11.5 and later, use `sp_countmetadata` to determine the current value of number of open objects and number of open databases. For example:

```
1> sp_countmetadata 'open objects'
2> go
```

There are 338 user objects in all database(s), requiring 165 Kbytes of memory. The 'open objects' configuration parameter is currently set to 500.

```
1> sp_countmetadata 'open databases'
2> go
```

There are 5 databases, requiring 189 Kbytes of memory. The 'open databases' configuration parameter is currently set to 12.

## os\_attach\_region errors

Message text                      `os_attach_region: shmat(%d): %s`

---

**Note** This error may be caused by a hardware problem.

---

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Explanation            | <p>Adaptive Server uses the function <code>os_attach_region</code> to attach to a shared memory region based on a shared memory identifier. This error occurs when Adaptive Server fails to attach to a region. In this message, (<code>%d</code>) is the shared memory identifier and (<code>%s</code>) is an operating system error message.</p> <hr/> <p><b>Note</b> The message is raised on UNIX systems only. Other operating systems raise slightly different errors.</p> <hr/> |
| Action                 | <p>An <code>os_attach_region</code> error can occur when the value of the operating system parameter <code>shmseg</code> is too small. For information about configuring shared memory properly, refer to the Adaptive Server Enterprise installation and configuration guides for your platform and version.</p> <p>This message can also be raised by Monitor Server. Refer to the <i>Monitor Server User's Guide</i> for details.</p>                                               |
| Additional information | Refer to the operating system man pages for the <code>shmat()</code> system call.                                                                                                                                                                                                                                                                                                                                                                                                      |

## os\_create\_region errors

|              |                                                                                                                                                                                                                                             |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Message text | <pre>os_create_region: shmget (0x%x): %s os_create_region: Shared memory segment %d is in the way os_create_region: uninitialized shared memory descriptor os_create_region: shmat (%d): %s os_create_region: can't allocate %d bytes</pre> |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

---

**Note** This error may be caused by a hardware problem.

---

|             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Explanation | <p>Adaptive Server uses the following functions to manage shared memory:</p> <ul style="list-style-type: none"> <li>• <code>os_get_shmid</code> – create a shared memory identifier.</li> <li>• <code>os_create_region</code> – create a region based on a shared memory identifier.</li> <li>• <code>os_attach_region</code> – attach to a region based on a shared memory identifier.</li> <li>• <code>os_detach_region</code> – detach from (and delete) the shared region.</li> <li>• <code>os_format_shmid</code> – format a shared memory identifier for printing.</li> </ul> |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

When the os\_create\_region errors occur, Adaptive Server will not start.

---

**Note** The message texts shown here apply to UNIX systems only. Other operating systems raise slightly different errors.

---

**os\_create\_region: shmget (0x<%x>): <%s>**

This message is written to the error log when Adaptive Server fails to get a shared memory segment. In this message, <%x> is a shared memory key based on the shared memory identifier and <%s> is an operating system error message.

**os\_create\_region: Shared memory segment <%d> is in the way**

This error follows the shmget message and is also written to the Adaptive Server error log. A value of -1 for <%d> means the region does not exist.

**os\_create\_region: Uninitialized shared memory descriptor**

During creation of a shared memory region, Adaptive Server attempts to validate the descriptor for the memory region. This message is written to the error log if the descriptor is found to be invalid.

**os\_create\_region: shmat (<%d>): <%s>**

This message is written to the error log when Adaptive Server fails to attach at an address. In this message, <%d> is the shared memory identifier and <%s> is an operating system error message.

**os\_create\_region: Can't allocate <%d> bytes**

Adaptive Server is unable to allocate the number of bytes it requested for the shared memory region.

Action

- 1 At the operating system level, check which shared memory processes are using Adaptive Server, and whether any shared memory segments are being used by Adaptive Server.

On UNIX, run this command as the “sybase” user:

```
% ipcs -m
IPC status from workstation1 as of Fri May 26 14:08:25 2007
T ID KEY MODE OWNER GROUP
Shared Memory:
m 256 0x699b7e24 --rw----- sybase sybase
m 257 0x699b7e25 --rw----- sybase sybase
```

If shared memory segments are being used by Adaptive Server, reboot the operating system to clear shared memory or remove them using the `ipcrm` operating system command.

---

**Note** Before removing the shared memory segments, identify the process that created them using the command `ipcs -ma` to ensure you only remove the appropriate segments.

---

- 2 Check the `$SYBASE` directory to determine whether there are any `*.krg` or `*.srg` files left from an abnormal Adaptive Server exit. If any such files exist, delete them.

`os_create_region` errors can occur when shared memory is not configured properly on your operating system. For information about configuring shared memory properly, refer to:

- The Adaptive Server installation and configuration guide for your platform and version
- *System Administration Guide: Volume 2*, Chapter 3, “Configuring Memory

### Shared memory error on Digital Unix

The error message “`os_create_region: can't allocate %d bytes`” indicates that one or more kernel parameters needs to be reset. Logically, resetting `shm-max` should allow Adaptive Server to configure shared memory. However, other operating system kernel parameters also affect allocation. Consult your operating system documentation for details.

Additional information

- Refer to the operating system man pages for the `shget()` and `shmat()` system calls.
- Refer to the operating system man pages for `ipcs` and `ipcrm`.

## Read/write errors

Message text

```
kernel: write error on virtual disk %d, block %ld
or:
```

```
kernel: read error on virtual disk %d, block %ld
```

---

**Note** This error may be caused by a hardware problem.

---

Explanation

This error occurs when Adaptive Server fails to read or write to the specified location on disk. This failure is usually a result of a physical disk problem.

If mirroring is enabled, you may also see these errors:

```
udunmirror: i/o error on %s device '%s'
```

or

```
sddone: %s error on virtual device
```

Action

If you see error 823 followed by `sddone` errors on an IBM Regatta server, the problem could be caused by a cache line synchronization problem on this family of servers. See *TechNote 1023360* for details and the solution from IBM at <http://www.sybase.com/detail?id=1023360>.

If you are running on another platform, use this procedure to correct the error:

- 1 Identify the device with the problem by selecting the row from `sysdevices` that has the virtual disk number indicated in the error message:

```
1> select name, phyname from master..sysdevices
2> where low/power(2,24) = <virtual_disk_number>
3> go
```

The output from this query should provide the physical name of the damaged disk.

---

**Note** The `sp_diskblock` procedure described below also identifies the device.

---

- 2 Shut down Adaptive Server and examine the disk as soon as feasible and correct any problems found to prevent further corruption, then restart Adaptive Server.
- 3 Use the procedure described below to determine which database (if there is more than one database on the device), table, or index is affected by this error. Use this information to assess the potential severity of the problem quickly and decide on an action that is appropriate for your operating environment.



- a Use the `sp_diskblock` procedure supplied in “Gathering information about read/write errors” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*. The syntax of this procedure is:

```
sp_diskblock <virtual_disk>, <block_number>
```

For example, if the read/write error message displays a virtual disk number 4 and a block number 871, the procedure command is:

```
1> sp_diskblock 4, 871
2> go
```

Virtual disk 4, block 871 corresponds to:

```
Logical page 1895 in the "production" database (dbid=4) on device
"main".
```

- b Use the `dbcc page` command to determine which object is using that page. Refer to “Finding an object name from a page number” in the chapter “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery* for details on this procedure.

---

**Warning!** Use the `dbcc page` command only as directed. This command is undocumented, nonstandard, and is provided “as is” without any warranty. Sybase Technical Support does not support this command for general use. Although the command syntax is provided for reference, use this command only in the specific diagnostic situations described, and with the specific syntax shown. Failure to do so could cause performance problems and database corruption.

---

Additional information      Refer to “dopen error” on page 667 and “udunmirror errors” on page 691 for information about dopen, sddone, and udunmirror errors.

## Server is unnamed

Message text                      server is unnamed

**Explanation** This is an informational message that displays during Adaptive Server start-up if the Adaptive Server has not been given a name using the `sp_addserver` system procedure. The Adaptive Server name is stored in the `syssservers` system table as well as in the `@@servername` global variable. This name is used primarily for managing server-to-server communication and does not have to match the server name used either in the `interfaces` file or in the `RUN_<server>` file.

**Action** No action is required. However, the Adaptive Server administrator can prevent this message from displaying by assigning a name to the Adaptive Server using the `sp_addserver` procedure. Refer to the *Reference Manual: Procedures* for information about restrictions on valid Adaptive Server names added with `sp_addserver`.

To assign a name to an Adaptive Server using the `sp_addserver` command (and eliminate the message), type:

```
1> sp_addserver <server_name>, local
2> go
```

Where `<server_name>` is the name you want to assign to the Adaptive Server.

The `@@servername` global variable will not reflect the name until Adaptive Server is restarted. All subsequent starts of the Adaptive Server display the new name during Adaptive Server start-up. For example, for an Adaptive Server named “MY\_SERVER”, the message would be:

```
server: server name is 'MY_SERVER'
```

Although this Adaptive Server name is associated with server-to-server communication, specifying a name does not in itself allow this communication and therefore does not represent a security risk. Naming all Adaptive Servers this way makes the Adaptive Server error logs easy to identify, because the Adaptive Server name is contained within the log.

## Stack guardword corrupted

**Message text** \*\*\* Stack guardword corrupted.

Explanation This error occurs when the stack area used by an Adaptive Server process is corrupted. Adaptive Server usually shuts itself down after this error to avoid problems that could arise from the stack corruption.

---

**Warning!** If this error occurs and Adaptive Server remains running, shut it down immediately!

---

At start-up, Adaptive Server allocates one stack area for every configured user connection. These stacks are in contiguous areas of memory, with a guard at the end of each stack. At the end of each stack guard area is a “guardword,” which is a 4-byte structure with a pattern. Adaptive Server periodically checks this pattern to determine whether it has changed. A change indicates that a process has overflowed its stack guard area. When this occurs, Adaptive Server prints the following messages in the error log and shuts down:

```
kernel: *** Stack overflow detected: <limit>: 0x%lx <sp>: 0x%lx
kernel: *** Stack guardword corrupted
kernel: *** Stack corrupted, server aborting
```

In the first message, *<limit>* is the address of the end of the stack guard area, and *<sp>* is the current value of the stack pointer.

In addition, Adaptive Server periodically checks the stack pointer to determine whether it is completely outside both the stack and the stack guard area for the process. If it is, Adaptive Server shuts down, even if the stack guardword is not corrupted, and prints the following messages in the error log:

```
kernel: *** Stack overflow detected: limit: 0x%lx sp: 0x%lx
kernel: *** Stack corrupted, server aborting
```

Stack corruption is usually the result of one process’s stack overflowing its boundaries and writing on top of another process’ stack. In general, an Adaptive Server process writes past the end of its stack area for one of two reasons:

- The stack size is too small. The information placed on the stack is valid but its total size exceeds the specified stack size. This is often caused by a single, complex query that requires more than the average amount of stack space.
- An incorrectly handled error results in recursive calls that eventually fill the stack, regardless of its size.

Action

If this error occurs because a complex query has run out of stack area, correct the error by using one of the following methods.

---

**Note** Choose solution 1, or choose between solutions 2 and 3 depending on your scenario.

---

**Solution 1**

Break up the query into smaller pieces, if possible.

**Solution 2 – Server remains running after error occurs**

Restart the server. Use the following information to select an appropriate size for the stack size configuration parameter:

- The stack size should always be an even increment of 2K (2048 bytes). If a stack size is specified that is not divisible by 2K, Adaptive Server rounds it up to the next 2K increment.
- Increase the stack size by no more than 2K at a time. If the stack size is too large, you may be unable to start Adaptive Server.
- Increasing the stack size results in Adaptive Server requiring more memory for each configured user connection and device.

Use `sp_configure` to increase the stack size configuration parameter:

```
1> sp_configure "stack size", <new_value>
2> go
```

Where `<new_value>` is the new stack size.

Once this change is made, shut down and restart Adaptive Server for the change to take effect.

---

**Warning!** If you increase the value of the stack size configuration parameter, you may need to adjust other memory-related parameters. Refer to “Memory Use and Performance” in *Performance and Tuning: Basics* and “Configuring Memory” in the *System Administration Guide: Volume 2*.

---

**Solution 3 – Server shuts down after error occurs**

- 1 Restart the Adaptive Server.
- 2 Increase the stack size, following the instructions in “Solution 2 – Server remains running after error occurs.”

- 3 Increase the stack guard size by 2K. Use `sp_configure` to increase this parameter:

```
1> sp_configure "stack guard size", <new_value>
2> go
```

- 4 Once this change is made, restart Adaptive Server to have the change take affect.

#### Additional information

Increasing the stack size results in Adaptive Server requiring more memory for each configured user connection and device. Therefore, increasing the stack size without increasing the memory parameter, will result in less total cache space for Adaptive Server and could affect performance.

Refer to “Setting Configuration Parameters” in the *System Administration Guide: Volume 1* for information about using the `sp_configure` procedure and the stack configuration parameters. Also see *Reference Manual: Procedures* for additional information on `sp_configure`.

## time slice error

#### Message text


```
timeslice %d
```

#### Explanation

The Adaptive Server Enterprise scheduler manages the processing time provided to client tasks. The scheduler allows a task to execute on a server engine using the rule:

$$\langle \text{time slice} \rangle + \langle \text{cpu grace time} \rangle = \langle \text{maximum continuous processing time} \rangle$$

The following diagram illustrates management of processing time using default values for time slice (1 clock tick, 1/10th of a second) and `cpu grace time` (500 clock ticks, 50 seconds). The task begins execution at the end of the “free time.”



At periodic yield points, the task checks to see if the time slice has been exceeded. If it has not, the task continues executing. If it has, the task voluntarily yields the engine within the `cpu grace time`; if the task had not yet fulfilled the client request, the task goes to the end of the run queue. If there are no other tasks in the run queue, the server grants the task another processing interval.

It is possible for a task not to encounter a yield point and to exceed the time slice period. When this occurs, the server terminates the task with a time slice error at the end of the `cpu grace time`.

time slice errors are typically due to a delay in acquiring an Adaptive Server resource, a coding problem, or some other issue; for example, the Server may have made a blocking system call that does not return for a long time.

Action

If you receive time slice errors, take these steps:

- 1 Set up a configurable shared memory dump (CSMD) using “time slice” as the dump condition. Information about CSMD can be found in the section “Other Useful Tasks” in the most recent version of *Troubleshooting and Disaster Recovery*.
- 2 Try increasing the value of the `cpu grace time` parameter. See “Configuration parameters” in the *System Administration Guide: Volume 1* chapter “Setting Configuration Parameters.”

---

**Warning!** Sybase does not recommend changing the time slice configuration parameter to resolve these errors, only the `cpu grace time` parameter. Changing the time slice parameter affects all tasks and changes the dynamics of Adaptive Server’s processing time management; changing the `cpu grace time` affects only tasks that would be terminated with a time slice error.

---

If `cpu grace time` is set to 1500 and time slice errors persist, contact Sybase Technical Support.

Additional information

If you contact Sybase Technical Support, have on hand the information that is listed in “Reporting errors” on page 4, including the output of the Adaptive Server error log and the output of the CSMD dump.

See “Scheduling Client Task Processing Time” in the most recent version of *Performance and Tuning: Basics* for additional information about the time slice parameter.

Version in which this error is raised

All versions

## t\_rcvconnect error

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Message text           | <code>nconnect: t_rcvconnect, %s</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|                        | <hr/> <p><b>Note</b> This error may be caused by a hardware problem.</p> <hr/>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Explanation            | <p>This error occurs when an Adaptive Server using a TLI-based interface tries to establish a client connection and an operating system error message is returned. This can happen when Adaptive Server tries to execute an RPC on a server that is not running (Adaptive Server, Open Server, or Backup Server).</p> <p>When this error occurs, the connection is not established and the message is written to the error log. (%s) is the text of the operating system error that caused the failure.</p> <p>This is not a serious error unless it occurs frequently.</p>                                                                                       |
| Action                 | <p>Resolve the problem that caused the operating system error contained in the Adaptive Server error message. You may need assistance from an operating system or network administrator. Refer to the operating system man pages for information about the <code>t_rcvconnect()</code> system call.</p> <p>Check the Adaptive Server error log. If other errors exist, refer to the write-ups in this manual for information about recovering from them.</p> <p>If, after resolving the operating system problem and the problems found in the error log, you still get the message “<code>nconnect: t_rcvconnect, %s,</code>” call Sybase Technical Support.</p> |
| Additional information | <p>Before calling Sybase Technical Support, have the information available that is listed in “Reporting errors” on page 4, including the contents of the <i>interfaces</i> file and the output of <code>sp_helpserver</code> and <code>sp_configure</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                   |

## udunmirror errors

|              |                                                                                                                                                                                           |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Message text | <pre>udunmirror: i/o error on primary device '%s' udunmirror: i/o error on secondary device '%s' udunmirror: failing over to '%s'</pre>                                                   |
| Explanation  | <p>This error occurs when Adaptive Server detects an I/O error on the primary or secondary device of a mirrored pair of devices. It is usually the result of a physical disk problem.</p> |

### I/O error detected on primary device

If an I/O error is detected on the primary device, Adaptive Server disables mirroring. `sysdevices` reflects information for the device as if a `disk unmirror` with `side = primary` and `mode = retain` was issued. Adaptive Server does this automatically and writes kernel messages like the following to the error log:

```
kernel: udunmirror: i/o error on primary device '<primary_device>'
kernel: DataServer i/o to the device will be disabled
kernel: udunmirror: failing over to 'secondary_device'
```

### I/O error detected on secondary device

If an I/O error is detected on the secondary device, Adaptive Server disables mirroring and `sysdevices` reflects information for the device as if a `disk unmirror` with `side = secondary` and `mode = retain` was issued. Adaptive Server does this automatically and writes kernel messages like the following to the error log:

```
kernel: udunmirror: i/o error on secondary device '/usr/u/sybase/test.mir'
kernel: DataServer i/o to the device will be disabled
```

#### Action

- 1 Using your operating system procedures, determine whether the device is damaged. If it is damaged, repair or replace it.

If the device is not damaged but it was unavailable for some other reason, make it available again.

- 2 If you use operating system files for Sybase devices or if the name of the device was changed, drop the mirror and re-create it:

```
1> disk unmirror name = "<device_name>",
2> side = {primary|secondary}, mode=remove
3> go

1> disk mirror name = "<device_name>",
2> mirror = "<full_mirror_name>"
3> go
```

Where `side` is the device that was temporarily disabled.

---

**Note** On a UNIX system, you can also use the operating system command `touch` to create the file, then issue a `disk remirror` command.

---

- 3 If you did not need to drop and re-create the mirror, use the `disk remirror` command to make the device accessible to Adaptive Server:

```
1> disk remirror name = "<device_name>"
```



```
2> go
```

---

**Note** Use the `disk remirror` command if Adaptive Server automatically disabled mirroring, or if you issued a `disk unmirror` with `mode = retain` command. If you issued a `disk unmirror` with `mode = remove` command, you can only reinstate disk mirroring by issuing the `disk mirror` command.

---

Additional information Refer to the *Reference Manual: Commands* for information about using the `disk mirror`, `disk remirror`, and `disk unmirror` commands.

## No more engines can be offlined

Message text `ueoffline: No more engines can be offlined in order to support %d outstanding network connections. Please try later when the server has fewer connections.`

Explanation Adaptive Server Enterprise version 12.0 and later allows System Administrators to bring down server engines using the `dbcc engine(offline)` command. Along with the existing `dbcc engine(online)` command, this makes it possible for System Administrators to dynamically optimize the number of engines on a server based on the work load.

When you execute `dbcc engine(offline)`, Adaptive Server checks whether the remaining engines can accommodate the outstanding user connections. This error is raised if taking the engine offline would make it impossible to support current connections to the server.

Action Wait until there are fewer active connections, then re-try the `dbcc engine(offline)` command.

Use `sp_who` periodically to determine the number of active user connections on your server. Use the following command to determine the upper limit on the number of file descriptors available for server connections:

```
1> select <@@max_connections>
2> go
```

The return value represents the maximum possible descriptors available. Take into account overhead such as site handler connections, and note that overhead increases with the number of engines.

Version in which this error is raised Version 12.0 and later

*No more engines can be offlined*

---