

SYBASE®

Reference Manual: Building Blocks

**Adaptive Server® Enterprise**

15.0

DOCUMENT ID: DC36271-01-1500-02

LAST REVISED: October 2005

Copyright © 1987-2005 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, the Sybase logo, ADA Workbench, Adaptable Windowing Environment, Adaptive Component Architecture, Adaptive Server, Adaptive Server Anywhere, Adaptive Server Enterprise, Adaptive Server Enterprise Monitor, Adaptive Server Enterprise Replication, Adaptive Server Everywhere, Adaptive Warehouse, Afaia, Answers Anywhere, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-Translator, APT-Library, AvantGo Mobile Delivery, AvantGo Mobile Inspection, AvantGo Mobile Marketing Channel, AvantGo Mobile Pharma, AvantGo Mobile Sales, AvantGo Pylon, AvantGo Pylon Application Server, AvantGo Pylon Conduit, AvantGo Pylon PIM Server, AvantGo Pylon Pro, Backup Server, BizTracker, ClearConnect, Client-Library, Client Services, Convoy/DM, Copernicus, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DataWindow .NET, DB-Library, dbQueue, Developers Workbench, DirectConnect, DirectConnect Anywhere, Distribution Director, e-ADK, E-Anywhere, e-Biz Impact, e-Biz Integrator, E-Whatever, EC Gateway, ECMAP, ECRTP, eFulfillment Accelerator, Embedded SQL, EMS, Enterprise Application Studio, Enterprise Client/Server, Enterprise Connect, Enterprise Data Studio, Enterprise Manager, Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, eProcurement Accelerator, EWA, Financial Fusion, Financial Fusion Server, Gateway Manager, GlobalFIX, iAnywhere, iAnywhere Solutions, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, M2M Anywhere, Mach Desktop, Mail Anywhere Studio, Mainframe Connect, Maintenance Express, Manage Anywhere Studio, M-Business Channel, M-Business Network, M-Business Server, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, mFolio, Mirror Activator, MySupport, Net-Gateway, Net-Library, New Era of Networks, ObjectConnect, ObjectCycle, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Biz, Open Client, Open Client/Connect, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, PB-Gen, PC APT Execute, PC DB-Net, PC Net Library, PocketBuilder, Pocket PowerBuilder, Power++, power.stop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, QAnywhere, Rapport, RemoteWare, RepConnector, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Report-Execute, Report Workbench, Resource Manager, RFID Anywhere, RW-DisplayLib, RW-Library, S-Designer, SDF, Search Anywhere, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILS, smart.partners, smart.parts, smart.script, SOA Anywhere, SQL Advantage, SQL Anywhere, SQL Anywhere Studio, SQL Code Checker, SQL Debug, SQL Edit, SQL Edit/TPU, SQL Everywhere, SQL Modeler, SQL Remote, SQL Server, SQL Server Manager, SQL SMART, SQL Toolset, SQL Server/CFT, SQL Server/DBM, SQL Server SNMP SubAgent, SQL Station, SQLJ, STEP, SupportNow, S.W.I.F.T. Message Format Libraries, Sybase Central, Sybase Client/Server Interfaces, Sybase Financial Server, Sybase Gateways, Sybase IQ, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SybFlex, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, TradeForce, Transact-SQL, Translation Toolkit, UltraLite, UltraLite.NET, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, VisualWriter, VQL, WarehouseArchitect, Warehouse Control Center, Warehouse Studio, Warehouse WORKS, Watcom, Watcom SQL, Watcom SQL Server, Web Deployment Kit, Web.PB, Web.SQL, WebSights, WebViewer, WorkGroup SQL Server, XA-Library, XA-Server, XcelleNet, and XP Server are trademarks of Sybase, Inc. 06/05

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

All other company and product names used herein may be trademarks or registered trademarks of their respective companies.

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

<b>About This Book .....</b>	<b>xi</b>	
<b>CHAPTER 1</b>	<b>System and User-Defined Datatypes .....</b>	<b>1</b>
	Datatype categories .....	1
	Range and storage size .....	2
	Datatypes of columns, variables, or parameters .....	4
	Declaring the datatype for a column in a table .....	5
	Declaring the datatype for a local variable in a batch or procedure	5
	5	
	Declaring the datatype for a parameter in a stored procedure..	5
	Determining the datatype of a literal.....	6
	Datatypes of mixed-mode expressions .....	7
	Determining the datatype hierarchy .....	7
	Determining precision and scale .....	9
	Datatype conversions.....	9
	Automatic conversion of fixed-length NULL columns .....	10
	Handling overflow and truncation errors.....	10
	Standards and compliance.....	11
	Exact numeric datatypes .....	12
	Integer types.....	13
	Decimal datatypes .....	14
	Standards and compliance .....	15
	Approximate numeric datatypes .....	16
	Understanding approximate numeric datatypes .....	16
	Range, precision, and storage size .....	17
	Entering approximate numeric data .....	17
	Values that may be entered by Open Client clients .....	17
	Standards and compliance .....	18
	Money datatypes .....	18
	Accuracy.....	18
	Range and storage size .....	18
	Entering monetary values.....	18
	Standards and compliance .....	19
	Timestamp datatype.....	19

- Creating a timestamp column..... 19
- Date and time datatypes ..... 20
  - Range and storage requirements..... 21
  - Entering date and time data ..... 21
  - Standards and compliance..... 25
- Character datatypes..... 25
  - unichar, univarchar..... 25
  - Length and storage size..... 26
  - Entering character data..... 27
  - Treatment of blanks..... 29
  - Manipulating character data..... 30
  - Standards and compliance..... 30
- Binary datatypes ..... 30
  - Valid binary and varbinary entries..... 30
  - Entries of more than the maximum column size ..... 31
  - Treatment of trailing zeros..... 31
  - Platform dependence ..... 32
  - Standards and compliance..... 33
- bit datatype..... 33
  - Standards and compliance..... 33
- sysname and longsysname datatypes ..... 33
  - Standards and compliance..... 34
- text, image, and unitext datatypes ..... 34
  - Data structures used for storing text, unitext, and image data 35
  - Initializing text, unitext, and image columns..... 36
  - Saving space by allowing NULL..... 37
  - Getting information from sysindexes ..... 37
  - Using readtext and writetext..... 38
  - Determining how much space a column uses..... 38
  - Restrictions on text, image, and unitext columns ..... 39
  - Selecting text, unitext, and image data ..... 39
  - Converting text and image datatypes..... 40
  - Converting to or from unitext..... 40
  - Pattern matching in text data..... 41
  - Duplicate rows..... 41
  - Standards and compliance..... 41
- User-defined datatypes ..... 41
  - Standards and compliance..... 42

- CHAPTER 2 Transact-SQL Functions ..... 43**
  - Types of functions ..... 43
  - Aggregate functions ..... 49
    - Aggregates used with group by..... 50
    - Aggregate functions and NULL values..... 50

Vector and scalar aggregates .....	50
Aggregate functions as row aggregates .....	53
Datatype conversion functions .....	55
Converting character data to a noncharacter type .....	58
Converting from one character type to another .....	58
Converting numbers to a character type .....	59
Rounding during conversion to and from money types .....	59
Converting date and time information .....	60
Converting between numeric types .....	60
Arithmetic overflow and divide-by-zero errors .....	61
Conversions between binary and integer types .....	62
Converting between binary and numeric or decimal types .....	63
Converting image columns to binary types .....	63
Converting other types to bit .....	63
Converting NULL value .....	64
Date functions .....	64
Date parts .....	64
Mathematical functions .....	65
Security functions .....	66
String functions .....	67
Limits on string functions .....	67
System functions .....	68
Text, unitext, and image columns .....	68
Text and image functions .....	69
abs .....	70
acos .....	71
ascii .....	72
asin .....	73
atan .....	74
atn2 .....	75
avg .....	76
audit_event_name .....	78
biginttohex .....	80
case .....	81
cast .....	84
ceiling .....	87
char .....	89
char_length .....	91
charindex .....	93
coalesce .....	94
col_length .....	96
col_name .....	97
compare .....	98
convert .....	103

cos.....	109
cot .....	110
count .....	111
count_big.....	113
current_date .....	115
current_time .....	116
curunreservedpgs .....	117
data_pages .....	119
datachange .....	121
datalength .....	123
dateadd .....	124
datediff .....	127
datename .....	130
datepart.....	132
day .....	136
db_id .....	137
db_name .....	138
degrees .....	139
derived_stat.....	140
difference .....	143
exp .....	144
floor .....	145
get_appcontext.....	147
getdate .....	149
getutcdate .....	150
has_role .....	151
hextobigint.....	153
hextoint.....	154
host_id.....	155
host_name .....	156
identity_burn_max.....	157
index_col .....	158
index_colorder.....	159
inttohex.....	160
is_quiesced .....	161
is_sec_service_on.....	163
isnull .....	164
lct_admin.....	165
left .....	168
len .....	170
license_enabled .....	171
list_appcontext .....	172
lockscheme .....	173
log .....	174

log10 .....	175
lower.....	176
ltrim .....	177
max .....	178
min .....	180
month .....	181
mut_excl_roles .....	182
newid.....	183
next_identity .....	185
nullif.....	186
object_id.....	188
object_name.....	189
pagesize .....	190
partition_id.....	192
partition_name .....	193
patindex.....	194
pi .....	197
power .....	198
proc_role .....	199
radians .....	201
rand .....	202
replicate.....	203
reserved_pages .....	204
reverse .....	206
right .....	207
rm_appcontext .....	209
role_contain.....	210
role_id .....	211
role_name .....	212
round.....	213
row_count.....	215
rtrim .....	216
set_appcontext.....	217
show_role.....	219
show_sec_services .....	220
sign.....	221
sin.....	222
sortkey.....	223
soundex.....	228
space.....	229
square .....	230
sqrt .....	231
str .....	232
str_replace .....	234

stuff .....	236
substring.....	238
sum .....	240
suser_id.....	242
suser_name .....	243
syb_quit.....	244
syb_sendmsg.....	245
tan .....	246
tempdb_id .....	247
textptr .....	248
textvalid .....	249
to_unichar .....	250
tran_dumptable_status.....	251
tsequal.....	252
uhighsurr .....	254
ulowsurr.....	255
upper .....	256
uscalar.....	257
used_pages.....	258
user .....	260
user_id .....	261
user_name .....	262
valid_name.....	263
valid_user.....	264
year .....	265
<b>CHAPTER 3</b>	
<b>Global Variables.....</b>	<b>267</b>
Adaptive Server global variables.....	267
<b>CHAPTER 4</b>	
<b>Expressions, Identifiers, and Wildcard Characters.....</b>	<b>275</b>
Expressions.....	275
Size of expressions .....	276
Arithmetic and character expressions .....	276
Relational and logical expressions.....	276
Operator precedence .....	277
Arithmetic operators .....	277
Bitwise operators.....	278
String concatenation operator .....	279
Comparison operators.....	280
Nonstandard operators.....	280
Using any, all and in.....	281
Negating and testing .....	281
Ranges .....	281



	Using nulls in expressions .....	281
	Connecting expressions .....	283
	Using parentheses in expressions .....	284
	Comparing character expressions.....	284
	Using the empty string.....	285
	Including quotation marks in character expressions .....	285
	Using the continuation character .....	285
	Identifiers.....	285
	Short identifiers .....	287
	Tables beginning with # (temporary tables) .....	288
	Case sensitivity and identifiers .....	288
	Uniqueness of object names .....	289
	Using delimited identifiers .....	289
	Identifying tables or columns by their qualified object name ..	290
	Determining whether an identifier is valid.....	292
	Renaming database objects.....	292
	Using multibyte character sets .....	292
	Pattern matching with wildcard characters.....	293
	Using not like .....	294
	Case and accent insensitivity .....	295
	Using wildcard characters .....	295
	Using multibyte wildcard characters.....	297
	Using wildcard characters as literal characters .....	297
	Using wildcard characters with datetime data .....	299
<b>CHAPTER 5</b>	<b>Reserved Words.....</b>	<b>301</b>
	Transact-SQL reserved words .....	301
	ANSI SQL reserved words .....	302
	Potential ANSI SQL reserved words .....	303
<b>CHAPTER 6</b>	<b>SQLSTATE Codes and Messages .....</b>	<b>305</b>
	Warnings .....	305
	Exceptions.....	306
	Cardinality violations .....	306
	Data exceptions.....	307
	Integrity constraint violations .....	308
	Invalid cursor states .....	308
	Syntax errors and access rule violations.....	309
	Transaction rollbacks .....	310
	with check option violation.....	310
<b>Index .....</b>		<b>313</b>



# About This Book

The *Adaptive Server Reference Manual* includes four guides to Sybase® Adaptive Server® Enterprise and the Transact-SQL® language:

- *Building Blocks* describes the “parts” of Transact-SQL: datatypes, built-in functions, global variables, expressions and identifiers, reserved words, and SQLSTATE errors. Before you can use Transact-SQL successfully, you must understand what these building blocks do and how they affect the results of Transact-SQL statements.
- *Commands* provides reference information about the Transact-SQL commands, which you use to create statements.
- *Procedures* provides reference information about system procedures, catalog stored procedures, extended stored procedures, and dbcc stored procedures. All procedures are created using Transact-SQL statements.
- *Tables* provides reference information about the system tables, which store information about your server, databases, users, and other details of your server. It also provides information about the tables in the dbccdb and dbccalt databases.

## Audience

The *Adaptive Server Reference Manual* is intended as a reference tool for Transact-SQL users of all levels.

## How to use this book

- Chapter 1, “System and User-Defined Datatypes,” describes the system and user-defined datatypes that are supplied with Adaptive Server and indicates how to use them to create user-defined datatypes.
- Chapter 2, “Transact-SQL Functions,” lists the Adaptive Server functions in a table that provides the name and a brief description.
- Chapter 3, “Global Variables,” lists the system-defined variables for Adaptive Server in a table that provides the name and a brief description of the returned status.
- Chapter 4, “Expressions, Identifiers, and Wildcard Characters,” which provides information about using the Transact-SQL language.

- 
- Chapter 5, “Reserved Words,” provides information about the Transact-SQL and ANSI SQL keywords.
  - Chapter 6, “SQLSTATE Codes and Messages,” contains information about Adaptive Server SQLSTATE status codes and the associated messages.

#### **Related documents**

The Adaptive Server<sup>®</sup> 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 *Installation Guide* for your platform – describes installation, upgrade, and configuration procedures for all Adaptive Server and related Sybase products.
- *What’s New in Adaptive Server Enterprise?* – describes the new features in Adaptive Server version 15.0, the system changes added to support those features, and changes that may affect your existing applications.
- *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.
- The *Configuration Guide* for your platform – provides instructions for performing specific configuration tasks for Adaptive Server.
- *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.
- *Historical Server User’s Guide* – describes how to use Historical Server to obtain performance information for SQL Server<sup>®</sup> and Adaptive Server.
- *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).
- *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 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.
- *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 and 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* – provides in-depth information about administering servers and databases. This manual includes 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.
  - *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.
  - *Unified Agent and Agent Management Console* – Describes the Unified Agent, which provides runtime services to manage, monitor and control distributed Sybase resources.
  - *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.
  - *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 in Adaptive Server Enterprise* – 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/>.

### 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 Select Products from the navigation bar on the left.
- 3 Select a product name from the product list and click Go.
- 4 Select the Certification Report filter, specify a time frame, and click Go.
- 5 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 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 EBFs/Maintenance releases is displayed.

Padlock icons indicate that you do not have download authorization for certain EBFs/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 EBFs/Maintenance report, or click the product description to download the software.

**Conventions**

The following sections describe conventions used in this manual.



SQL is a free-form language. There are no rules about the number of words you can put on a line or where you must break a line. However, for readability, all examples and most syntax statements in this manual are formatted so that each clause of a statement begins on a new line. Clauses that have more than one part extend to additional lines, which are indented. Complex commands are formatted using modified Backus Naur Form (BNF) notation.

Table 1 shows the conventions for syntax statements that appear in this manual:

**Table 1: Font and syntax conventions for this manual**

Element	Example
Command names, procedure names, utility names, and other keywords display in sans serif font.	<code>select</code> <code>sp_configure</code>
Database names and datatypes are in sans serif font.	<code>master database</code>
Book names, file names, variables, and path names are in italics.	<i>System Administration Guide</i> <i>sql.ini</i> file <i>column_name</i> \$SYBASE/ASE directory
Variables—or words that stand for values that you fill in—when they are part of a query or statement, are in italics in Courier font.	<code>select column_name</code> <code>from table_name</code> <code>where search_conditions</code>
Type parentheses as part of the command.	<code>compute row_aggregate (column_name)</code>
Double colon, equals sign indicates that the syntax is written in BNF notation. Do not type this symbol. Indicates “is defined as”.	<code>::=</code>
Curly braces mean that you must choose at least one of the enclosed options. Do not type the braces.	<code>{cash, check, credit}</code>
Brackets mean that to choose one or more of the enclosed options is optional. Do not type the brackets.	<code>[cash   check   credit]</code>
The comma means you may choose as many of the options shown as you want. Separate your choices with commas as part of the command.	<code>cash, check, credit</code>
The pipe or vertical bar ( ) means you may select only one of the options shown.	<code>cash   check   credit</code>
An ellipsis (...) means that you can <i>repeat</i> the last unit as many times as you like.	<code>buy thing = price [cash   check   credit]</code> <code>[, thing = price [cash   check   credit]]...</code> You must buy at least one thing and give its price. You may choose a method of payment: one of the items enclosed in square brackets. You may also choose to buy additional things: as many of them as you like. For each thing you buy, give its name, its price, and (optionally) a method of payment.

- Syntax statements (displaying the syntax and all options for a command) appear as follows:

```
sp_dropdevice [device_name]
```

For a command with more options:

```
select column_name
from table_name
where search_conditions
```

In syntax statements, keywords (commands) are in normal font and identifiers are in lowercase. Italic font shows user-supplied words.

- Examples showing the use of Transact-SQL commands are printed like this:

```
select * from publishers
```

- Examples of output from the computer appear 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)

In this manual, most of the examples are in lowercase. However, you can disregard case when typing Transact-SQL keywords. For example, SELECT, Select, and select are the same.

Adaptive Server's sensitivity to the case of database objects, such as table names, depends on the sort order installed on Adaptive Server. You can change case sensitivity for single-byte character sets by reconfiguring the Adaptive Server sort order. For more information, see the *System Administration Guide*.

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.

## Accessibility features

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.



# System and User-Defined Datatypes

This chapter describes the Transact-SQL datatypes, which specify the type, size, and storage format of columns, stored procedure parameters, and local variables.

<b>Topics</b>	<b>Page</b>
Datatype categories	1
Range and storage size	2
Datatypes of columns, variables, or parameters	4
Datatypes of mixed-mode expressions	7
Datatype conversions	9
Standards and compliance	11
Exact numeric datatypes	12
Approximate numeric datatypes	16
Money datatypes	18
Timestamp datatype	19
Date and time datatypes	20
Character datatypes	25
Binary datatypes	30
bit datatype	33
sysname and longsysname datatypes	33
text, image, and unitext datatypes	34
User-defined datatypes	41

## Datatype categories

Adaptive Server provides several system datatypes and the user-defined datatypes timestamp, sysname, and longsysname. Table 1-1 lists the categories of Adaptive Server datatypes. Each category is described in a section of this chapter.

**Table 1-1: Datatype categories**

Category	Used for
Exact numeric datatypes	Numeric values (both integers and numbers with a decimal portion) that must be represented exactly
Approximate numeric datatypes	Numeric data that can tolerate rounding during arithmetic operations
Money datatypes	Monetary data
Timestamp datatype	Tables that are browsed in Client-Library™ applications
Date and time datatypes	Date and time information
Character datatypes	Strings consisting of letters, numbers, and symbols
Binary datatypes	Raw binary data, such as pictures, in a hexadecimal-like notation
bit datatype	True/false and yes/no type data
sysname and longsysname datatypes	System tables
text, image, and unitext datatypes	Printable characters or hexadecimal-like data that requires more than the maximum column size provided by your server's logical page size.
Abstract datatypes	Adaptive Server supports abstract datatypes through Java classes. See <i>Java in Adaptive Server Enterprise</i> for more information.
User-defined datatypes	Defining objects that inherit the rules, default, null type, IDENTITY property, and base datatype of the datatypes listed in this table. text undergoes character-set conversion if client is using a different character set, image does not.

## Range and storage size

Table 1-2 lists the system-supplied datatypes and their synonyms and provides information about the range of valid values and storage size for each. For simplicity, the datatypes are printed in lowercase characters, although Adaptive Server allows you to use either uppercase or lowercase characters for system datatypes. User-defined datatypes, such as timestamp, are *case-sensitive*. Most Adaptive Server-supplied datatypes are not reserved words and can be used to name other objects.

**Table 1-2: Adaptive Server system datatypes**

Datatypes by category	Synonyms	Range	Bytes of storage
<i>Exact numeric: integers</i>			

Datatypes by category	Synonyms	Range	Bytes of storage
bigint		Whole numbers between $2^{63}$ and $-2^{63}$ - 1 (from -9,223,372,036,854,775,808 to +9,223,372,036,854,775,807, inclusive.	8
int	integer	$2^{31}$ - 1 (2,147,483,647) to $-2^{31}$ (-2,147,483,648)	4
smallint		$2^{15}$ - 1 (32,767) to $-2^{15}$ (-32,768)	2
tinyint		0 to 255 (Negative numbers are not permitted)	1
unsigned bigint		Whole numbers between 0 and 18,446,744,073,709,551,615	8
unsigned int		Whole numbers between 0 and 4,294,967,295	4
unsigned smallint		Whole numbers between 0 and 65535	2
<i>Exact numeric: decimals</i>			
numeric (p, s)		$10^{38}$ - 1 to $-10^{38}$	2 to 17
decimal (p, s)	dec	$10^{38}$ - 1 to $-10^{38}$	2 to 17
<i>Approximate numeric</i>			
float (precision)		machine dependent	4 for default precision < 16, 8 for default precision >= 16
double precision		machine dependent	8
real		machine dependent	4
<i>Money</i>			
smallmoney		214,748.3647 to -214,748.3648	4
money		922,337,203,685,477.5807 to -922,337,203,685,477.5808	8
<i>Date/time</i>			
smalldatetime		January 1, 1900 to June 6, 2079	4
datetime		January 1, 1753 to December 31, 9999	8
date		January 1, 0001 to December 31, 9999	4
time		12:00:00AM to 11:59:59:999PM	4
<i>Character</i>			
char(n)	character	pagesize	n

Datatypes by category	Synonyms	Range	Bytes of storage
varchar(n)	character varying, char varying	pagesize	actual entry length
unichar	Unicode character	pagesize	$n * @@unicharsize$ ( $@@unicharsize$ equals 2)
univarchar	Unicode character varying, char varying	pagesize	actual number of characters * $@@unicharsize$
nchar(n)	national character, national char	pagesize	$n * @@ncharsize$
nvarchar(n)	nchar varying, national char varying, national character varying	pagesize	$@@ncharsize * \text{number of}$ characters
text		$2^{31} - 1$ (2,147,483,647) bytes or fewer	0 when uninitialized; multiple of 2K after initialization
unitext		1 – 1,073,741,823	0 when uninitialized; multiple of 2K after initialization
<i>Binary</i>			
binary(n)		pagesize	$n$
varbinary(n)		pagesize	actual entry length
image		$2^{31} - 1$ (2,147,483,647) bytes or fewer	0 when uninitialized; multiple of 2K after initialization
<i>Bit</i>			
bit		0 or 1	1 (one byte holds up to 8 bit columns)

## Datatypes of columns, variables, or parameters

You must declare the datatype for a column, local variable, or parameter. The datatype can be any of the system-supplied datatypes, or any user-defined datatype in the database.



## Declaring the datatype for a column in a table

To declare the datatype of a new column in a create table or alter table statement, use:

```
create table [[database.]owner.]table_name
    (column_name datatype [identity | not null | null]
    [, column_name datatype [identity | not null |
    null]]...)

alter table [[database.]owner.]table_name
    add column_name datatype [identity | null]
    [, column_name datatype [identity | null]]...
```

For example:

```
create table sales_daily
    (stor_id char(4) not null,
    ord_num numeric(10,0) identity,
    ord_amt money null)
```

You can also declare the datatype of a new column in a select into statement, use `convert` or `cast`:

```
select convert(double precision, x), cast (int, y) into
    newtable from oldtable
```

## Declaring the datatype for a local variable in a batch or procedure

To declare the datatype for a local variable in a batch or stored procedure, use:

```
declare @variable_name datatype
    [, @variable_name datatype ]...
```

For example:

```
declare @hope money
```

## Declaring the datatype for a parameter in a stored procedure

Use the following syntax to declare the datatype for a parameter in a stored procedure:

```
create procedure [owner.]procedure_name [;number]
    [[(@parameter_name datatype [= default] [output]
    [, @parameter_name datatype [= default]
    [output]]...)]
    [with recompile]
as SQL_statements
```

For example:

```
create procedure auname_sp @auname varchar(40)
as
select au_lname, title, au_ord
from authors, titles, titleauthor
where @auname = au_lname
and authors.au_id = titleauthor.au_id
and titles.title_id = titleauthor.title_id
```

## Determining the datatype of a literal

### Numeric literals

Numeric literals entered with E notation are treated as float; all others are treated as exact numerics:

- Literals between  $2^{31} - 1$  and  $-2^{31}$  with no decimal point are treated as integer.
- Literals that include a decimal point, or that fall outside the range for integers, are treated as numeric.

---

**Note** To preserve backward compatibility, use E notation for numeric literals that should be treated as float.

---

### Character literals

In versions of Adaptive Server earlier than 12.5.1, when the client's character set was different from the server's character set, conversions were generally enabled to allow the text of SQL queries to be converted to the server's character set before being processed. If any character could not be converted because it could not be represented in the server's character set, the entire query was rejected. This character set "bottleneck" has been removed as of Adaptive Server version 12.5.1.

You cannot declare the datatype of a character literal. Adaptive Server treats character literals as `varchar`, except those that contain characters that cannot be converted to the server's default character set. Such literals are treated as `univarchar`. This makes it possible to perform such queries as selecting `unichar` data in a server configured for "iso\_1" using a "sjis" (Japanese) client. For example:

```
select * from mytable where unichar_column = ' 五 '
```

Since the character literal cannot be represented using the `char` datatype (in "iso\_1"), it is promoted to the `unichar` datatype, and the query succeeds.

## Datatypes of mixed-mode expressions

When you perform concatenation or mixed-mode arithmetic on values with different datatypes, Adaptive Server must determine the datatype, length, and precision of the result.

## Determining the datatype hierarchy

Each system datatype has a **datatype hierarchy**, which is stored in the `systypes` system table. User-defined datatypes inherit the hierarchy of the system datatype on which they are based.

The following query ranks the datatypes in a database by hierarchy. In addition to the information shown below, your query results will include information about any user-defined datatypes in the database:

```
select name, hierarchy
       from systypes
       order by hierarchy

name                                     hierarchy
-----
floatn                                   1
float                                     2
datetimn                                  3
datetime                                  4
real                                       5
numericn                                  6
numeric                                   7
```

decimaln	8
decimal	9
moneyn	10
money	11
smallmoney	12
smalldatetime	13
intn	14
uintn	15
bigint	16
ubigint	17
int	18
uint	19
smallint	20
usmallint	21
tinyint	22
bit	23
univarchar	24
unichar	25
unitext	26
sysname	27
varchar	27
nvarchar	27
longsysname	27
char	28
nchar	28
timestamp	29
varbinary	29
binary	30
text	31
image	32
date	33
time	34
datetime	35
datetime	36
extended type	99

---

**Note** `u<int type>` is an internal representation. The correct syntax for unsigned types is `unsigned {int | integer | bigint | smallint }`

---

The datatype hierarchy determines the results of computations using values of different datatypes. The result value is assigned the datatype that is closest to the top of the list or has the least hierarchical value.

In the following example, *qty* from the sales table is multiplied by *royalty* from the roysched table. *qty* is a `smallint`, which has a hierarchy of 20; *royalty* is an `int`, which has a hierarchy of 18. Therefore, the datatype of the result is an `int`:

```
smallint (qty) * int (royalty) = int
```

## Determining precision and scale

For numeric and decimal datatypes, each combination of precision and scale is a distinct Adaptive Server datatype. If you perform arithmetic on two numeric or decimal values:

- *n1* with precision *p1* and scale *s1*, and
- *n2* with precision *p2* and scale *s2*

Adaptive Server determines the precision and scale of the results as shown in Table 1-3.

**Table 1-3: Precision and scale after arithmetic operations**

Operation	Precision	Scale
$n1 + n2$	$\max(s1, s2) + \max(p1 - s1, p2 - s2) + 1$	$\max(s1, s2)$
$n1 - n2$	$\max(s1, s2) + \max(p1 - s1, p2 - s2) + 1$	$\max(s1, s2)$
$n1 * n2$	$s1 + s2 + (p1 - s1) + (p2 - s2) + 1$	$s1 + s2$
$n1 / n2$	$\max(s1 + p2 + 1, 6) + p1 - s1 + p2$	$\max(s1 + p2 - s2 + 1, 6)$

## Datatype conversions

Many conversions from one datatype to another are handled automatically by Adaptive Server. These are called implicit conversions. Other conversions must be performed explicitly with the `convert`, `hextoint`, `inttohex`, `hextobigint`, and `biginttohex` functions. See “Datatype conversion functions” on page 55 for details about datatype conversions supported by Adaptive Server.

## Automatic conversion of fixed-length NULL columns

Only columns with variable-length datatypes can store null values. When you create a NULL column with a fixed-length datatype, Adaptive Server automatically converts it to the corresponding variable-length datatype. Adaptive Server does not inform the user of the datatype change.

Table 1-4 lists the fixed- and variable-length datatypes to which they are converted. Certain variable-length datatypes, such as `money`, are reserved datatypes; you cannot use them to create columns, variables, or parameters:

**Table 1-4: Automatic conversion of fixed-length datatypes**

Original fixed-length datatype	Converted to
<code>char</code>	<code>varchar</code>
<code>unichar</code>	<code>univarchar</code>
<code>nchar</code>	<code>nvarchar</code>
<code>binary</code>	<code>varbinary</code>
<code>datetime</code>	<code>datetime</code>
<code>date</code>	<code>datetime</code>
<code>time</code>	<code>time</code>
<code>float</code>	<code>float</code>
<code>bigint</code> , <code>int</code> , <code>smallint</code> , and <code>tinyint</code>	<code>int</code>
<code>unsigned bigint</code> , <code>unsigned int</code> , and <code>unsigned smallint</code>	<code>int</code>
<code>decimal</code>	<code>decimal</code>
<code>numeric</code>	<code>numeric</code>
<code>money</code> and <code>smallmoney</code>	<code>money</code>

## Handling overflow and truncation errors

The `arithabort` option determines how Adaptive Server behaves when an arithmetic error occurs. The two `arithabort` options, `arithabort arith_overflow` and `arithabort numeric_truncation`, handle different types of arithmetic errors. You can set each option independently, or set both options with a single `set arithabort on` or `set arithabort off` statement.

- `arithabort arith_overflow` specifies behavior following a divide-by-zero error or a loss of precision during either an explicit or an implicit datatype conversion. This type of error is considered serious. The default setting, `arithabort arith_overflow on`, rolls back the entire transaction in which the error occurs. If the error occurs in a batch that does not contain a transaction, `arithabort arith_overflow on` does not roll back earlier commands in the batch, but Adaptive Server does not execute any statements that follow the error-generating statement in the batch.

Setting `arith_overflow` to `on` refers to the execution time, not to the level of normalization to which Adaptive Server is set.

If you set `arithabort arith_overflow off`, Adaptive Server aborts the statement that causes the error, but continues to process other statements in the transaction or batch.

- `arithabort numeric_truncation` specifies behavior following a loss of scale by an exact numeric datatype during an implicit datatype conversion. (When an explicit conversion results in a loss of scale, the results are truncated without warning.) The default setting, `arithabort numeric_truncation on`, aborts the statement that causes the error but continues to process other statements in the transaction or batch. If you set `arithabort numeric_truncation off`, Adaptive Server truncates the query results and continues processing.

The `arithignore` option determines whether Adaptive Server prints a warning message after an overflow error. By default, the `arithignore` option is turned off. This causes Adaptive Server to display a warning message after any query that results in numeric overflow. To ignore overflow errors, use `set arithignore on`.

## Standards and compliance

Table 1-5 lists the ANSI SQL standards and compliance levels for Transact-SQL datatypes.

**Table 1-5: ANSI SQL standards and compliance levels for Transact-SQL datatypes**

Transact-SQL – ANSI SQL datatypes	Transact-SQL extensions – User-defined datatypes
<ul style="list-style-type: none"> <li>• char</li> <li>• varchar</li> <li>• smallint</li> <li>• int</li> <li>• bigint</li> <li>• decimal</li> <li>• numeric</li> <li>• float</li> <li>• real</li> <li>• date</li> <li>• time</li> <li>• double precision</li> </ul>	<ul style="list-style-type: none"> <li>• binary</li> <li>• varbinary</li> <li>• bit</li> <li>• nchar</li> <li>• datetime</li> <li>• smalldatetime</li> <li>• tinyint</li> <li>• unsigned smallint</li> <li>• unsigned int</li> <li>• unsigned bigint</li> <li>• money</li> <li>• smallmoney</li> <li>• text</li> <li>• unitext</li> <li>• image</li> <li>• nvarchar</li> <li>• unichar</li> <li>• univarchar</li> <li>• sysname</li> <li>• longsysname</li> <li>• timestamp</li> </ul>

## Exact numeric datatypes

Use the exact numeric datatypes when you must represent a value exactly. Adaptive Server provides exact numeric types for both integers (whole numbers) and numbers with a decimal portion.



## Integer types

Adaptive Server provides the following exact numeric datatypes to store integers: bigint, int (or integer), smallint, tinyint and each of their unsigned counterparts. Choose the integer type based on the expected size of the numbers to be stored. Internal storage size varies by type, as shown in Table 1-6.

**Table 1-6: Integer datatypes**

Datatype	Stores	Bytes of storage
bigint	Whole numbers between $-2^{63}$ and $2^{63} - 1$ (from -9,223,372,036,854,775,808 to +9,223,372,036,854,775,807, inclusive).	8
int[eger]	Whole numbers between $-2^{31}$ and $2^{31} - 1$ (-2,147,483,648 and 2,147,483,647), inclusive.	4
smallint	Whole numbers between $-2^{15}$ and $2^{15} - 1$ (-32,768 and 32,767), inclusive.	2
tinyint	Whole numbers between 0 and 255, inclusive. (Negative numbers are not permitted.)	1
unsigned bigint	Whole numbers between 0 and 18,446,744,073,709,551,615	8
unsigned int	Whole numbers between 0 and 4,294,967,295	4
unsigned smallint	Whole numbers between 0 and 65,535	2

### Entering integer data

Enter integer data as a string of digits without commas. Integer data can include a decimal point as long as all digits to the right of the decimal point are zeros. The smallint, integer, and bigint datatypes can be preceded by an optional plus or minus sign. The tinyint datatype can be preceded by an optional plus sign.

Table 1-7 shows some valid entries for a column with a datatype of integer and indicates how isql displays these values:

**Table 1-7: Valid integer values**

Value entered	Value displayed
2	2
+2	2
-2	-2
2.	2
2.000	2

Table 1-8 lists some invalid entries for an integer column:

**Table 1-8: Invalid integer values**

Value entered	Type of error
2,000	Commas not allowed.
2-	Minus sign should precede digits.
3.45	Digits to the right of the decimal point are nonzero digits.

## Decimal datatypes

Adaptive Server provides two other exact numeric datatypes, numeric and dec[imal], for numbers that include decimal points. The numeric and decimal datatypes are identical in all respects but one: only numeric datatypes with a scale of 0 and integer datatypes can be used for the IDENTITY column.

### Specifying precision and scale

The numeric and decimal datatypes accept two optional parameters, precision and scale, enclosed in parentheses and separated by a comma:

*datatype* [(*precision* [, *scale*])]

Adaptive Server treats each combination of precision and scale as a distinct datatype. For example, numeric(10,0) and numeric(5,0) are two separate datatypes. The precision and scale determine the range of values that can be stored in a decimal or numeric column:

- The precision specifies the maximum number of decimal digits that can be stored in the column. It includes *all* digits, both to the right and to the left of the decimal point. You can specify precisions ranging from 1 digit to 38 digits or use the default precision of 18 digits.
- The scale specifies the maximum number of digits that can be stored to the right of the decimal point. The scale must be less than or equal to the precision. You can specify a scale ranging from 0 digits to 38 digits, or use the default scale of 0 digits.

### Storage size

The storage size for a numeric or decimal column depends on its precision. The minimum storage requirement is 2 bytes for a 1- or 2-digit column. Storage size increases by approximately 1 byte for each additional 2 digits of precision, up to a maximum of 17 bytes.

Use the following formula to calculate the exact storage size for a numeric or decimal column:

$\text{ceiling}(\text{precision} / \log_{10}(256)) + 1$

For example, the storage size for a numeric(18,4) column is 9 bytes.

Entering decimal data Enter decimal and numeric data as a string of digits preceded by an optional plus or minus sign and including an optional decimal point. If the value exceeds either the precision or scale specified for the column, Adaptive Server returns an error message. Exact numeric types with a scale of 0 are displayed without a decimal point.

Table 1-9 shows some valid entries for a column with a datatype of numeric(5,3) and indicates how these values are displayed by isql:

**Table 1-9: Valid decimal values**

Value entered	Value displayed
12.345	12.345
+12.345	12.345
-12.345	-12.345
12.345000	12.345
12.1	12.100
12	12.000

Table 1-10 shows some invalid entries for a column with a datatype of numeric(5,3):

**Table 1-10: Invalid decimal values**

Value entered	Type of error
1,200	Commas not allowed.
12-	Minus sign should precede digits.
12.345678	Too many nonzero digits to the right of the decimal point.

## Standards and compliance

Transact-SQL provides the smallint, int, bigint, numeric, and decimal ANSI SQL exact numeric datatypes. The unsigned bigint, unsigned int, unsigned smallint, and tinyint type is a Transact-SQL extension.

## Approximate numeric datatypes

Use the approximate numeric types, float, double precision, and real, for numeric data that can tolerate rounding. The approximate numeric types are especially suited to data that covers a wide range of values. They support all aggregate functions and all arithmetic operations except modulo.

### Understanding approximate numeric datatypes

Approximate numeric datatypes, used to store floating-point numbers, are inherently slightly inaccurate in their representation of real numbers—hence the name “approximate numeric.” To use these datatypes, you must understand their limitations.

When a floating-point number is printed or displayed, the printed representation is not quite the same as the stored number, and the stored number is not quite the same as the number that the user entered. Most of the time, the stored representation is close enough, and software makes the printed output look just like the original input, but you must understand the inaccuracy if you plan to use floating-point numbers for calculations, particularly if you are doing repeated calculations using approximate numeric datatypes—the results can be surprisingly and unexpectedly inaccurate.

The inaccuracy occurs because floating-point numbers are stored in the computer as binary fractions (that is, as a representative number divided by a power of 2), but the numbers we use are decimal (powers of 10). This means that only a very small set of numbers can be stored accurately: 0.75 ( $3/4$ ) can be stored accurately because it is a binary fraction (4 is a power of 2); 0.2 ( $2/10$ ) cannot (10 is not a power of 2).

Some numbers contain too many digits to store accurately. double precision is stored as 8 binary bytes and can represent about 17 digits with reasonable accuracy. real is stored as 4 binary bytes and can represent only about 6 digits with reasonable accuracy.

If you begin with numbers that are almost correct, and perform computations with them using other numbers that are almost correct, you can easily end up with a result that is not even close to being correct. If these considerations are important to your application, use an exact numeric datatype.

## Range, precision, and storage size

The real and double precision types are built on types supplied by the operating system. The float type accepts an optional binary precision in parentheses. float columns with a precision of 1–15 are stored as real; those with higher precision are stored as double precision.

The range and storage precision for all three types is machine-dependent.

Table 1-11 shows the range and storage size for each approximate numeric type. isql displays only 6 significant digits after the decimal point and rounds the remainder:

**Table 1-11: Approximate numeric datatypes**

Datatype	Bytes of storage
float[(default precision)]	4 for default precision < 16 8 for default precision >= 16
double precision	8
real	4

## Entering approximate numeric data

Enter approximate numeric data as a mantissa followed by an optional exponent:

- The mantissa is a signed or unsigned number, with or without a decimal point. The column's binary precision determines the maximum number of binary digits allowed in the mantissa.
- The exponent, which begins with the character “e” or “E,” must be a whole number.

The value represented by the entry is the following product:

$$\text{mantissa} * 10^{\text{EXPONENT}}$$

For example, 2.4E3 represents the value 2.4 times  $10^3$ , or 2400.

## Values that may be entered by Open Client clients

“NaN” and “Inf” are special values that the floating point number standard uses to represent values that are “not a number” and “infinity,” respectively. Adaptive Server does not usually permit these values, but Open Client clients, particularly native-mode bcp, can force these values into tables.

## Standards and compliance

ANSI SQL – Compliance level: The float, double precision, and real datatypes are entry-level compliant.

## Money datatypes

Use the money and smallmoney datatypes to store monetary data. You can use these types for U.S. dollars and other decimal currencies, but Adaptive Server provides no means to convert from one currency to another. You can use all arithmetic operations except modulo, and all aggregate functions, with money and smallmoney data.

## Accuracy

Both money and smallmoney are accurate to one ten-thousandth of a monetary unit, but they round values up to two decimal places for display purposes. The default print format places a comma after every three digits.

## Range and storage size

Table 1-12 summarizes the range and storage requirements for money datatypes:

**Table 1-12: Money datatypes**

Datatype	Range	Bytes of storage
money	Monetary values between +922,337,203,685,477.5807 and -922,337,203,685,477.5808	8
smallmoney	Monetary values between +214,748.3647 and -214,748.3648	4

## Entering monetary values

Monetary values entered with E notation are interpreted as float. This may cause an entry to be rejected or to lose some of its precision when it is stored as a money or smallmoney value.

money and smallmoney values can be entered with or without a preceding currency symbol, such as the dollar sign (\$), yen sign (¥), or pound sterling sign (£). To enter a negative value, place the minus sign after the currency symbol. Do not include commas in your entry.

## Standards and compliance

ANSI SQL – The money and smallmoney datatypes are Transact-SQL extensions.

## Timestamp datatype

Use the user-defined timestamp datatype in tables that are to be browsed in Client-Library™ applications (see “Browse Mode” for more information). Adaptive Server updates the timestamp column each time its row is modified. A table can have only one column of timestamp datatype.

## Creating a *timestamp* column

If you create a column named timestamp without specifying a datatype, Adaptive Server defines the column as a timestamp datatype:

```
create table testing
(c1 int, timestamp, c2 int)
```

You can also explicitly assign the timestamp datatype to a column named timestamp:

```
create table testing
(c1 int, timestamp timestamp, c2 int)
```

or to a column with another name:

```
create table testing
(c1 int, t_stamp timestamp, c2 int)
```

You can create a column named timestamp and assign it another datatype (although this may be confusing to other users and does not allow the use of the browse functions in Open Client™ or with the tsequal function):

```
create table testing
```

```
(c1 int, timestamp datetime)
```

## Date and time datatypes

Use `datetime`, `smalldatetime`, `date`, and `time` to store absolute date and time information. Use `timestamp` to store binary-type information.

Adaptive Server has various ways to identify date and time. In versions earlier than 12.5.1, only `datetime` and `smalldatetime` were available. As of version 12.5.1, `date` and `time` are these separate datatypes:

- `date`
- `time`
- `smalldatetime`
- `datetime`

The default display format for dates is “Apr 15 1987 10:23PM”. You can use the `convert` function for other styles of date display. You can also perform some arithmetic calculations on date and time values with the built-in date functions, though Adaptive Server may round or truncate millisecond values.

- `datetime` columns hold dates between January 1, 1753 and December 31, 9999. `datetime` values are accurate to 1/300 second on platforms that support this level of granularity. Storage size is 8 bytes: 4 bytes for the number of days since the base date of January 1, 1900 and 4 bytes for the time of day.
- `smalldatetime` columns hold dates from January 1, 1900 to June 6, 2079, with accuracy to the minute. Its storage size is 4 bytes: 2 bytes for the number of days after January 1, 1900, and 2 bytes for the number of minutes after midnight.
- `date` columns hold dates from January 1, 0001 to December 31, 9999. Storage size is 4 bytes.
- `time` is between 00:00:00:000 and 23:59:59:999. You can use either military time or 12AM for noon and 12PM for midnight. A time value must contain either a colon or the AM or PM signifier. AM or PM may be in either uppercase or lowercase.

When entering date and time information, always enclose the time or date in single or double quotes.



## Range and storage requirements

Table 1-13 summarizes the range and storage requirements for the `datetime`, `smalldatetime`, `date`, and `time` datatypes:

**Table 1-13: Transact-SQL datatypes for storing dates and times**

Datatype	Range	Bytes of storage
<code>datetime</code>	January 1, 1753 through December 31, 9999	8
<code>smalldatetime</code>	January 1, 1900 through June 6, 2079	4
<code>date</code>	January 1, 0001 to December 31, 9999	4
<code>time</code>	12:00:00 AM to 11:59:59:999 PM	4

## Entering date and time data

The `datetime` and `smalldatetime` datatypes consist of a date portion either followed by or preceded by a time portion. (You can omit either the date or the time, or both.) The `date` datatype has only a date and the `time` datatype has only the time. You must enclose values in single or double quotes.

### Entering the date

Dates consist of a month, day, and year and can be entered in a variety of formats for `date`, `datetime`, and `smalldatetime`:

- You can enter the entire date as an unseparated string of 4, 6, or 8 digits, or use slash (/), hyphen (-), or period (.) separators between the date parts.
  - When entering dates as unseparated strings, use the appropriate format for that string length. Use leading zeros for single-digit years, months, and days. Dates entered in the wrong format may be misinterpreted or result in errors.
  - When entering dates with separators, use the `set dateformat` option to determine the expected order of date parts. If the first date part in a separated string is four digits, Adaptive Server interprets the string as `yyyy-mm-dd` format.
- Some date formats accept 2-digit years (`yy`):
  - Numbers less than 50 are interpreted as 20`yy`. For example, 01 is 2001, 32 is 2032, and 49 is 2049.
  - Numbers equal to or greater than 50 are interpreted as 19`yy`. For example, 50 is 1950, 74 is 1974, and 99 is 1999.

- You can specify the month as either a number or a name. Month names and their abbreviations are language-specific and can be entered in uppercase, lowercase, or mixed case.
- If you omit the date portion of a datetime or smalldatetime value, Adaptive Server uses the default date of January 1, 1900.

Table 1-14 describes the acceptable formats for entering the date portion of a datetime or smalldatetime value:

**Table 1-14: Date formats for date and time datatypes**

Date format	Interpretation	Sample entries	Meaning
4-digit string with no separators	Interpreted as yyyy. Date defaults to Jan 1 of the specified year.	“1947”	Jan 1 1947
6-digit string with no separators	Interpreted as <i>yyymmdd</i> . For yy < 50, year is 20yy. For yy >= 50, year is 19yy.	“450128” “520128”	Jan 28 2045 Jan 28 1952
8-digit string with no separators	Interpreted as <i>yyyymmdd</i> .	“19940415”	Apr 15 1994
String consisting of 2-digit month, 2-digit day, and year separated by slashes, hyphens, or periods, or a combination of the above	The <i>dateformat</i> and language set options determine the expected order of date parts. For <i>us_english</i> , the default order is <i>mdy</i> .  For yy < 50, year is interpreted as 20yy. For yy >= 50, year is interpreted as 19yy.	“4/15/94” “4.15.94” “4-15-94” “04.15/94”	All of these entries are interpreted as Apr 15 1994 when the <i>dateformat</i> option is set to <i>mdy</i> .
String consisting of 2-digit month, 2-digit day, and 4-digit year separated by slashes, hyphens, or periods, or a combination of the above	The <i>dateformat</i> and language set options determine the expected order of date parts. For <i>us_english</i> , the default order is <i>mdy</i> .	“04/15.1994”	Interpreted as Apr 15 1994 when the <i>dateformat</i> option is set to <i>mdy</i> .
Month is entered in character form (either full month name or its standard abbreviation), followed by an optional comma	If 4-digit year is entered, date parts can be entered in any order.  If day is omitted, all 4 digits of year must be specified. Day defaults to the first day of the month.  If year is only 2 digits (yy), it is expected to appear after the day. For yy < 50, year is interpreted as 20yy. For yy >= 50, year is interpreted as 19yy.	“April 15, 1994” “1994 15 apr” “1994 April 15” “15 APR 1994”  “apr 1994”  “mar 16 17” “apr 15 94”	All of these entries are interpreted as Apr 15 1994.  Apr 1 1994  Mar 16 2017 Apr 15 1994
The empty string “”	Date defaults to Jan 1 1900.	“”	Jan 1 1900

Entering the time

The time component of a datetime, smalldatetime, or time value must be specified as follows:

`hours[:minutes[:seconds[:milliseconds]]] [AM | PM]`

- Use 12AM for midnight and 12PM for noon.
- A time value must contain either a colon or an AM or PM signifier. The AM or PM can be entered in uppercase, lowercase, or mixed case.
- The seconds specification can include either a decimal portion preceded by a decimal point, or a number of milliseconds preceded by a colon. For example, “15:30:20:1” means twenty seconds and one millisecond past 3:30 PM; “15:30:20.1” means twenty and one-tenth of a second past 3:30 PM.
- If you omit the time portion of a datetime or smalldatetime value, Adaptive Server uses the default time of 12:00:00:000AM.

Displaying formats for *datetime*, *smalldatetime*, and *date* values

The display format for datetime and smalldatetime values is “Mon dd yyyy hh:mmAM” (or “PM”); for example, “Apr 15 1988 10:23PM”. To display seconds and milliseconds, and to obtain additional date styles and date-part orders, use the convert function to convert the data to a character string. Adaptive Server may round or truncate millisecond values.

Table 1-15 lists some examples of datetime entries and their display values:

**Table 1-15: Examples of datetime and date entries**

Entry	Value displayed
“1947”	Jan 1 1947 12:00AM
“450128 12:30:1PM”	Jan 28 2045 12:30PM
“12:30.1PM 450128”	Jan 28 2045 12:30PM
“14:30.22”	Jan 1 1900 2:30PM
“4am”	Jan 1 1900 4:00AM
<i>Examples of date</i>	
“1947”	Jan 1 1947
“450128”	Jan 28 2045
“520317”	Mar 17 1952

Displaying formats for *time* value

The display format for time values is “hh:mm:ss:mmmAM” (or “PM”); for example, “10:23:40:022PM”.

**Table 1-16: Examples of time entries**

Entry	Value displayed
"12:12:00"	12:12PM
"01:23PM" or "01:23:1PM"	1:23PM
"02:24:00:001"	2:24AM

Finding values that match a pattern

Use the like keyword to look for dates that match a particular pattern. If you use the equality operator (=) to search date or time values for a particular month, day, and year, Adaptive Server returns only those values for which the time is precisely 12:00:00:000AM.

For example, if you insert the value "9:20" into a column named arrival\_time, Adaptive Server converts the entry into "Jan 1 1900 9:20AM." If you look for this entry using the equality operator, it is not found:

```
where arrival_time = "9:20" /* does not match */
```

You can find the entry using the like operator:

```
where arrival_time like "%9:20%"
```

When using like, Adaptive Server first converts the dates to datetime or date format and then to varchar. The display format consists of the 3-character month in the current language, 2 characters for the day, 4 characters for the year, the time in hours and minutes, and "AM" or "PM."

When searching with like, you cannot use the wide variety of input formats that are available for entering the date portion of datetime, smalldatetime, date, and time values. Since the standard display formats do not include seconds or milliseconds, you cannot search for seconds or milliseconds with like and a match pattern, unless you are also using *style* 9 or 109 and the convert function.

If you are using like, and the day of the month is a number between 1 and 9, insert 2 spaces between the month and the day to match the varchar conversion of the datetime value. Similarly, if the hour is less than 10, the conversion places 2 spaces between the year and the hour. The following clause with 1 space between "May" and "2") finds all dates from May 20 through May 29, but not May 2:

```
like "May 2%"
```

You do not need to insert the extra space with other date comparisons, only with like, since the datetime values are converted to varchar only for the like comparison.

Manipulating dates

You can do some arithmetic calculations on date and time datatypes values with the built-in date functions. See "Date functions" on page 64.

## Standards and compliance

ANSI SQL – Compliance level: The datetime and smalldatetime datatypes are Transact-SQL extensions. date and time datatypes are entry-level compliant.

## Character datatypes

Which datatype you use for a situation depends on the type of data you are storing:

- Use the character datatypes to store strings consisting of letters, numbers, and symbols.
- Use `varchar(n)` and `char(n)` for both single-byte character sets such as `us_english` and for multibyte character sets such as Japanese.
- Use the `unichar(n)` and `univarchar(n)` datatypes to store Unicode characters. They are useful for single-byte or multibyte characters when you need a fixed number of bytes per character.
- Use the fixed-length datatype, `nchar(n)`, and the variable-length datatype, `nvarchar(n)`, for both single-byte and multibyte character sets, such as Japanese. The difference between `nchar(n)` and `char(n)` and `nvarchar(n)` and `varchar(n)` is that both `nchar(n)` and `nvarchar(n)` allocate storage based on  $n$  times the number of bytes per character (based on the default character set). `char(n)` and `varchar(n)` allocate  $n$  bytes of storage.
- Character datatypes can store a maximum of a page size worth of data
- Use the text datatype (described in “text, image, and unitext datatypes” on page 34)—or multiple rows in a subtable—for strings longer than the `char` or `varchar` datatype allow.

### ***unichar, univarchar***

You can use the `unichar` and `univarchar` datatypes anywhere that you can use `char` and `varchar` character datatypes, without having to make syntax changes.

In Adaptive Server version 12.5.1 and later, queries containing character literals that cannot be represented in the server's character set are automatically promoted to the unichar datatype so you do not have to make syntax changes for data manipulation language (DML) statements. Additional syntax is available for specifying arbitrary characters in character literals, but the decision to “promote” a literal to unichar is based solely on representability.

With data definition language (DDL) statements, the syntax changes required are minimal. For example, in the create table command, the size of a Unicode column is specified in units of 16-bit Unicode values, not bytes, thereby maintaining the similarity between char(200) and unichar(200). sp\_help, which reports on the lengths of columns, uses the same units. The multiplication factor (2) is stored in the new global variable @@unicharsize.

See Chapter 8, “Configuring Character Sets, Sort Orders, and Languages,” in the *System Administration Guide* for more information about Unicode.

## Length and storage size

Character variables strip the trailing spaces from strings when the variable is populated in a varchar column of a cursor.

Use  $n$  to specify the number of bytes of storage for char and varchar datatypes. For unichar, use  $n$  to specify the number of Unicode characters (the amount of storage allocated is 2 bytes per character). For nchar and nvarchar,  $n$  is the number of characters (the amount of storage allocated is  $n$  times the number of bytes per character for the server's current default character set).

If you do not use  $n$  to specify the length:

- The default length is 1 byte for columns created with create table, alter table, and variables created with declare.
- The default length is 30 bytes for values created with the convert function.

Entries shorter than the assigned length are blank-padded; entries longer than the assigned length are truncated without warning, unless the string\_truncation option to the set command is set to on. Fixed-length columns that allow nulls are internally converted to variable-length columns.

Use  $n$  to specify the maximum length in characters for the variable-length datatypes, `varchar(n)`, `univarchar(n)`, and `nvarchar(n)`. Data in variable-length columns is stripped of trailing blanks; storage size is the actual length of the data entered. Data in variable-length variables and parameters retains all trailing blanks, but is not padded to the defined length. Character literals are treated as variable-length datatypes.

Fixed-length columns tend to take more storage space than variable-length columns, but are accessed somewhat faster. Table 1-17 summarizes the storage requirements of the different character datatypes:

**Table 1-17: Character datatypes**

Datatype	Stores	Bytes of storage
<code>char(n)</code>	Character	$n$
<code>unichar(n)</code>	Unicode character	$n * @@unicharsize$ ( $@@unicharsize$ equals 2)
<code>nchar(n)</code>	National character	$n * @@ncharsize$
<code>varchar(n)</code>	Character varying	Actual number of characters entered
<code>univarchar(n)</code>	Unicode character varying	Actual number of characters * $@@unicharsize$
<code>nvarchar(n)</code>	National character varying	Actual number of characters * $@@ncharsize$

Determining column length with system functions

Use the `char_length` string function and `datalength` system function to determine column length:

- `char_length` returns the number of characters in the column, stripping trailing blanks for variable-length datatypes.
- `datalength` returns the number of bytes, stripping trailing blanks for data stored in variable-length columns.

When a `char` value is declared to allow NULL values, Adaptive Server stores it internally as a `varchar`.

If the `min` or `max` aggregate functions are used on a `char` column, the result returned is `varchar`, and is therefore stripped of all trailing spaces.

## Entering character data

Character strings must be enclosed in single or double quotes. If you use `set quoted_identifier on`, use single quotes for character strings; otherwise, Adaptive Server treats them as identifiers.

Strings that include the double-quote character should be surrounded by single quotes. Strings that include the single-quote character should be surrounded by double quotes. For example:

```
'George said, "There must be a better way."'  
"Isn't there a better way?"
```

An alternative is to enter two quotation marks for each quotation mark you want to include in the string. For example:

```
"George said, ""There must be a better way.""  
'Isn''t there a better way?'
```

To continue a character string onto the next line of your screen, enter a backslash (\) before going to the next line.

For more information about quoted identifiers, see the section “Delimited identifiers” of the *Transact SQL User’s Guide*.

## Entering Unicode characters

Optional syntax allows you to specify arbitrary Unicode characters. If a character literal is immediately preceded by U& or u& (with no intervening white space), the parser recognizes escape sequences within the literal. An escape sequence of the form \xxxx (where xxxx represents four hexadecimal digits) is replaced with the Unicode character whose scalar value is xxxx. Similarly, an escape sequence of the form \+yyyyyy is replaced with the Unicode character whose scalar value is yyyyyy. The escape sequence \\ is replaced by a single \. For example, the following is equivalent to:

```
select * from mytable where unichar_column = ' 五 '
```

```
select * from mytable where unichar_column = U&'\4e94'
```

The U& or u& prefix simply enables the recognition of escapes. The datatype of the literal is chosen solely on the basis of representability. Thus, for example, the following two queries are equivalent:

```
select * from mytable where char_column = 'A'
```

```
select * from mytable where char_column = U&'\0041'
```

In both cases, the datatype of the character literal is char, since “A” is an ASCII character, and ASCII is a subset of all Sybase-supported server character sets.

The U& and u& prefixes also work with the double-quoted character literals and for quoted identifiers. However, quoted identifiers must be representable in the server’s character set, insofar as all database objects are identified by names in system tables, and all such names are of datatype char.



## Treatment of blanks

The following example creates a table named `spaces` that has both fixed- and variable-length character columns:

```
create table spaces (cnot char(5) not null,
                   cnull char(5) null,
                   vnot varchar(5) not null,
                   vnull varchar(5) null,
                   explanation varchar(25) not null)

insert spaces values ("a", "b", "c", "d",
                    "pads char-not-null only")
insert spaces values ("1   ", "2   ", "3   ",
                    "4   ", "truncates trailing blanks")
insert spaces values ("  e", "  f", "  g",
                    "  h", "leading blanks, no change")
insert spaces values ("  w ", "  x ", "  y ",
                    "  z ", "truncates trailing blanks")
insert spaces values ("", "", "", "",
                    "empty string equals space ")

select "[" + cnot + "]",
       "[" + cnull + "]",
       "[" + vnot + "]",
       "[" + vnull + "]",
       explanation from spaces
       explanation
```

```
-----
[a   ] [b]   [c]   [d]   pads char-not-null only
[1   ] [2]   [3]   [4]   truncates trailing blanks
[  e] [  f] [  g] [  h] leading blanks, no change
[  w ] [  x] [  y] [  z] truncates trailing blanks
[    ] [ ]   [ ]   [ ]   empty string equals space
```

(5 rows affected)

This example illustrates how the column's datatype and null type interact to determine how blank spaces are treated:

- Only `char` not null and `nchar` not null columns are padded to the full width of the column; `char` null columns are treated like `varchar` and `nchar` null columns are treated like `nvarchar`.
- Only `unichar` not null columns are padded to the full width of the column; `unichar` null columns are treated like `univarchar`.
- Preceding blanks are not affected.

- Trailing blanks are truncated except for char, unichar, and nchar not null columns.
- The empty string (“ ”) is treated as a single space. In char, nchar, and unichar not null columns, the result is a column-length field of spaces.

## Manipulating character data

You can use the like keyword to search character strings for particular characters and the built-in string functions to manipulate their contents. You can use strings consisting of numbers for arithmetic after being converted to exact and approximate numeric datatypes with the convert function.

## Standards and compliance

ANSI SQL – Compliance level: Transact-SQL provides the char and varchar ANSI SQL datatypes. The nchar, nvarchar, unichar, and univarchar datatypes are Transact-SQL extensions.

## Binary datatypes

Use the binary datatypes, binary(*n*) and varbinary(*n*), to store raw binary data, such as pictures, in a raw binary notation, up to the maximum column size for your server’s logical page size.

## Valid *binary* and *varbinary* entries

Binary data begins with the characters “0x” and can include any combination of digits, and the uppercase and lowercase letters A through F.

Use *n* to specify the column length in bytes, or use the default length of 1 byte. Each byte stores 2 binary digits. If you enter a value longer than *n*, Adaptive Server truncates the entry to the specified length without warning or error.

Use the fixed-length binary type, binary(*n*), for data in which all entries are expected to be approximately equal in length.

Use the variable-length binary type, `varbinary(n)`, for data that is expected to vary greatly in length.

Because entries in binary columns are zero-padded to the column length ( $n$ ), they may require more storage space than those in `varbinary` columns, but they are accessed somewhat faster.

If you do not use  $n$  to specify the length:

- The default length is 1 byte for columns created with `create table`, `alter table`, and variables created with `declare`.
- The default length is 30 bytes for values created with the `convert` function.

## Entries of more than the maximum column size

Use the `image` datatype to store larger blocks of binary data (up to 2,147,483,647 bytes) on external data pages. You cannot use the `image` datatype for variables or for parameters in stored procedures. For more information, see “text, image, and unitext datatypes” on page 34.

## Treatment of trailing zeros

All binary not null columns are padded with zeros to the full width of the column. Trailing zeros are truncated in all `varbinary` data and in binary null columns, since columns that accept null values must be treated as variable-length columns.

The following example creates a table with all four variations of binary and `varbinary` datatypes, `NULL`, and `NOT NULL`. The same data is inserted in all four columns and is padded or truncated according to the datatype of the column.

```
create table zeros (bnot binary(5) not null,
                  bnull binary(5) null,
                  vnot varbinary(5) not null,
                  vnull varbinary(5) null)

insert zeros values (0x12345000, 0x12345000, 0x12345000, 0x12345000)
insert zeros values (0x123, 0x123, 0x123, 0x123)

select * from zeros
```

bnot	bnull	vnot	vnull
-----	-----	-----	-----
0x1234500000	0x123450	0x123450	0x123450
0x0123000000	0x0123	0x0123	0x0123

Because each byte of storage holds 2 binary digits, Adaptive Server expects binary entries to consist of the characters “0x” followed by an even number of digits. When the “0x” is followed by an odd number of digits, Adaptive Server assumes that you omitted the leading 0 and adds it for you.

Input values “0x00” and “0x0” are stored as “0x00” in variable-length binary columns (binary null, image, and varbinary columns). In fixed-length binary (binary not null) columns, the value is padded with zeros to the full length of the field:

```
insert zeros values (0x0, 0x0,0x0, 0x0)
select * from zeros where bnot = 0x00
bnot          bnull          vnot          vnull
-----
0x0000000000  0x00          0x00          0x00
```

If the input value does not include the “0x”, Adaptive Server assumes that the value is an ASCII value and converts it. For example:

```
create table sample (col_a binary(8))

insert sample values ('002710000000aeb1b')

select * from sample
col_a
-----
0x3030323731303030
```

## Platform dependence

The exact form in which you enter a particular value depends upon the platform you are using. Therefore, calculations involving binary data can produce different results on different machines.

You cannot use the aggregate functions sum or avg with the binary datatypes.

For platform-independent conversions between hexadecimal strings and integers, use the intohex and hextoint functions rather than the platform-specific convert function. For details, see “Datatype conversion functions” on page 55.

## Standards and compliance

ANSI SQL – Compliance level: The binary and varbinary datatypes are Transact-SQL extensions.

## *bit* datatype

Use the bit datatype for columns that contain true/false and yes/no types of data. The status column in the syscolumns system table indicates the unique offset position for bit datatype columns.

bit columns hold either 0 or 1. Integer values other than 0 or 1 are accepted, but are always interpreted as 1.

Storage size is 1 byte. Multiple bit datatypes in a table are collected into bytes. For example, 7 bit columns fit into 1 byte; 9 bit columns take 2 bytes.

Columns with a datatype of bit cannot be NULL and cannot have indexes on them.

## Standards and compliance

ANSI SQL – Compliance level: Transact-SQL extension.

## *sysname* and *longsysname* datatypes

*sysname* and *longsysname* are user-defined datatypes that are distributed on the Adaptive Server installation tape and used in the system tables. The definitions are:

- *sysname* – varchar(30) "not null"
- *longsysname* – varchar(255) "not null"

You can declare a column, parameter, or variable to be of types *sysname* and *longsysname*. Alternately, you can also create a user-defined datatype with a base type of *sysname* and *longsysname*, and then define columns, parameters, and variables with the user-defined datatype.

## Standards and compliance

ANSI SQL – Compliance level: All user-defined datatypes, including `sysname` and `longsysname`, are Transact-SQL extensions.

## text, image, and unitext datatypes

text columns are variable-length columns that can hold up to 2,147,483,647 ( $2^{31} - 1$ ) bytes of printable characters.

The variable-length unitext datatype can hold up to 1,073,741,823 Unicode characters (2,147,483,646 bytes).

image columns are variable-length columns that can hold up to 2,147,483,647 ( $2^{31} - 1$ ) bytes of raw binary data.

A key distinction between text and image is that text is subject to character-set conversion if you are not using the default character set of Adaptive Server default. image is not subject to character-set conversion.

Define a text, unitext, or image column as you would any other column, with a create table or alter table statement. text, unitext, or image datatype definitions do not include lengths. text, unitext, and image columns do permit null values. Their column definition takes the form:

```
column_name {text | image | unitext} [null]
```

For example, the create table statement for the author's blurbs table in the pubs2 database with a text column, blurb, that permits null values, is:

```
create table blurbs
  (au_id id not null,
  copy text null)
```

This example creates a unitext column that allows null values:

```
create table tb (ut unitext null)
```

To create the au\_pix table in the pubs2 database with an image column:

```
create table au_pix
  (au_id          char(11) not null,
  pic            image null,
  format_type    char(11) null,
  bytesize      int null,
  pixwidth_hor  char(14) null,
  pixwidth_vert char(14) null)
```

Adaptive Server stores text, *untext*, and image data in a linked list of data pages that are separate from the rest of the table. Each text, *untext*, or image page stores one logical page size worth of data (2, 4, 8, or 16K). All text, *untext*, and image data for a table is stored in a single page chain, regardless of the number of text, *untext*, and image columns the table contains.

You can place subsequent allocations for text, *untext*, and image data pages on a different logical device with `sp_placeobject`.

image values that have an odd number of hexadecimal digits are padded with a leading zero (an insert of “0xaaabb” becomes “0x0aaabb”).

You can use the `partition` option of the `alter table` command to partition a table that contains text, *untext*, and image columns. Partitioning the table creates additional page chains for the other columns in the table, but has *no* effect on the way the text, *untext*, and image columns are stored.

You can use *untext* anywhere you use the text datatype, with the same semantics. *untext* columns are stored in UTF-16 encoding, regardless of the Adaptive Server default character set.

## Data structures used for storing *text*, *untext*, and *image* data

When you allocate text, *untext*, or image data, a 16-byte text pointer is inserted into the row you allocated. Part of this text pointer refers to a text page number at the head of the text, *untext*, or image data. This text pointer is known as the first text page.

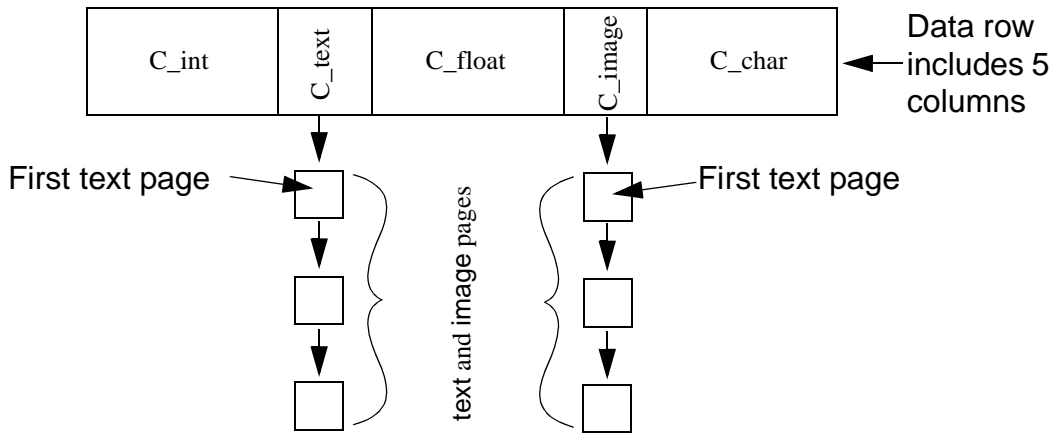
The first text page contains two parts:

- The text data page chain, which contains the text and image data and is a double-linked list of text pages
- The optional text-node structure, which is used to access the user text data

Once an first text page is allocated for text, *untext*, or image data, it is never deallocated. If an update to an existing text, *untext*, or image data row results in fewer text pages than are currently allocated for this text, *untext*, or image data, Adaptive Server deallocates the extra text pages. If an update to text, *untext*, or image data sets the value to NULL, all pages except the first text page are deallocated.

Figure 1-1 shows the relationship between the data row and the text pages.

Figure 1-1: Relationship between the text pointer and data rows



In Figure 1-1, columns `c_text` and `c_image` are text and image columns containing the pages at the bottom of the picture.

## Initializing *text*, *unitext*, and *image* columns

`text`, `unitext`, and `image` columns are not initialized until you update them or insert a non-null value. Initialization allocates at least one data page for each non-null text, `unitext`, or image data value. It also creates a pointer in the table to the location of the text, `unitext`, or image data.

For example, the following statements create the table `testtext` and initialize the `blurb` column by inserting a non-null value. The column now has a valid text pointer, and the first text page has been allocated.

```
create table testtext
(title_id varchar(6), blurb text null, pub_id char(4))

insert testtext values
("BU7832", "Straight Talk About Computers is an
annotated analysis of what computers can do for you: a
no-hype guide for the critical user.", "1389")
```

The following statements create a table for image values and initialize the image column:

```
create table imagetest
(image_id varchar(6), imagecol image null, graphic_id
char(4))
```



```
insert imagetest values
("94732", 0x0000008300000000000100000000013c, "1389")
```

---

**Note** Surround text values with quotation marks and precede image values with the characters “0x”.

---

For information on inserting and updating text, unitext, and image data with Client-Library programs, see the *Client-Library/C Reference Manual*.

## Defining unitext columns

You can define a unitext column the same way you define other datatypes, using create table or alter table statements. You do not define the length of a unitext column, and the column can be null.

This example creates a unitext column that allows null values:

```
create table tb (ut unitext null)
```

default unicode sort order defines the sort order for unitext columns for pattern matching in like clauses and in the patindex function, this is independent of the Adaptive Server default sort order.

## Saving space by allowing NULL

To save storage space for empty text, unitext, or image columns, define them to permit null values and insert nulls until you use the column. Inserting a null value does not initialize a text, unitext, or image column and, therefore, does not create a text pointer or allocate storage. For example, the following statement inserts values into the title\_id and pub\_id columns of the testtext table created above, but does not initialize the blurb text column:

```
insert testtext
(title_id, pub_id) values ("BU7832", "1389")
```

## Getting information from *sysindexes*

Each table with text, unitext, or image columns has an additional row in sysindexes that provides information about these columns. The name column in sysindexes uses the form “tablename.” The indid is always 255. These columns provide information about text storage:

**Table 1-18: Storage of text and image data**

Column	Description
ioampg	Pointer to the allocation page for the text page chain
first	Pointer to the first page of text data
root	Pointer to the last page
segment	Number of the segment where the object resides

You can query the `sysindexes` table for information about these columns. For example, the following query reports the number of data pages used by the `blurbs` table in the `pubs2` database:

```
select name, data_pages(db_id(), object_id("blurbs"), indid)
   from sysindexes
  where name = "tblurbs"
```

---

**Note** The system tables poster shows a one-to-one relationship between `sysindexes` and `systabstats`. This is correct, except for text and image columns, for which information is not kept in `systabstats`.

---

## Using *readtext* and *writetext*

Before you can use `writetext` to enter text data or `readtext` to read it, you must initialize the text column. For details, see `readtext` and `writetext` in *Reference Manual: Commands*.

Using `update` to replace existing *text*, *unitext*, and *image* data with `NULL` reclaims all allocated data pages except the first page, which remains available for future use of `writetext`. To deallocate all storage for the row, use `delete` to remove the entire row.

There are restrictions for using `readtext` and `writetext` on a column defined for *unitext*. For more information see the “Usage” sections under `readtext` and `writetext` in the *Reference Manual: Commands*.

## Determining how much space a column uses

`sp_spaceused` provides information about the space used for text data as `index_size`:

```
sp_spaceused blurbs
```

name	rowtotal	reserved	data	index_size	unused
-----	-----	-----	-----	-----	-----
blurbs	6	32 KB	2 KB	14 KB	16 KB

## Restrictions on *text*, *image*, and *unitext* columns

You cannot use *text*, *image*, or *unitext* columns:

- As parameters to stored procedures or as values passed to these parameters
- As local variables
- In order by clause, compute clause, group by, and union clauses
- In an index
- In subqueries or joins
- In a where clause, except with the keyword like
- With the + concatenation operator

## Selecting *text*, *unitext*, and *image* data

The following global variables return information on *text*, *unitext*, and *image* data:

**Table 1-19: *text*, *unitext*, and *image* global variables**

Variable	Explanation
<code>@@textptr</code>	The text pointer of the last <i>text</i> , <i>unitext</i> , or <i>image</i> column inserted or updated by a process. Do not confuse this global variable with the <code>textptr</code> function.
<code>@@textcolid</code>	ID of the column referenced by <code>@@textptr</code> .
<code>@@textdbid</code>	ID of a database containing the object with the column referenced by <code>@@textptr</code> .
<code>@@textobjid</code>	ID of the object containing the column referenced by <code>@@textptr</code> .
<code>@@textsize</code>	Current value of the <code>set textsize</code> option, which specifies the maximum length, in bytes, of <i>text</i> , <i>unitext</i> , or <i>image</i> data to be returned with a <code>select</code> statement. It defaults to 32K. The maximum size for <code>@@textsize</code> is $2^{31} - 1$ (that is, 2,147,483,647).
<code>@@textts</code>	Text timestamp of the column referenced by <code>@@textptr</code> .

## Converting *text* and *image* datatypes

You can explicitly convert text values to char, unichar, varchar, and univarchar, and image values to binary or varbinary with the convert function, but you are limited to the maximum length of the character and binary datatypes, which is determined by the maximum column size for your server’s logical page size. If you do not specify the length, the converted value has a default length of 30 bytes. Implicit conversion is not supported.

## Converting to or from unitext

You can implicitly convert any character or binary datatype to unitext, as well as explicitly convert to and from unitext to other datatypes. The conversion result, however, is limited to the maximum length of the destination datatype. When a unitext value cannot fit the destination buffer on a Unicode character boundary, data is truncated. If you have enabled enable surrogate processing, the unitext value is never truncated in the middle of a surrogate pair of values, which means that fewer bytes may be returned after the datatype conversion. For example, if a unitext column ut in table tb stores the string “U+0041U+0042U+00c2” (U+0041 representing the Unicode character “A”), this query returns the value “AB” if the server’s character set is UTF-8, because U+00C2 is converted to 2-byte UTF-8 0xc382:

```
select convert(char(3), ut) from tb
```

**Table 1-20: Converting to and from unitext**

These datatypes convert implicitly <b>to</b> unitext	These datatypes convert implicitly <b>from</b> unitext	These datatypes convert explicitly <b>from</b> unitext
char, varchar, unichar, univarchar, binary, varbinary, text, image	text, image	char, varchar, unichar, univarchar, binary, varbinary

The alter table modify command does not support text, image, or unitext columns to be the modified column. To migrate from a text to a unitext column:

- Use bcp out -Jutf8 out to copy text column data out
- Create a table with unitext columns
- Use bcp in -Jutf8 to insert data into the new table

## Pattern matching in *text* data

Use the `patindex` function to search for the starting position of the first occurrence of a specified pattern in a `text`, `unitext`, `varchar`, `univarchar`, `unichar`, or `char` column. The `%` wildcard character must precede and follow the pattern (except when you are searching for the first or last character).

You can also use the `like` keyword to search for a particular pattern. The following example selects each `text` data value from the `copy` column of the `blurbs` table that contains the pattern “Net Etiquette.”

```
select copy from blurbs
where copy like "%Net Etiquette%"
```

## Duplicate rows

The pointer to the `text`, `image`, and `unitext` data uniquely identifies each row. Therefore, a table that contains `text`, `image`, and `unitext` data does not contain duplicate rows unless there are rows in which all `text`, `image`, and `unitext` data is `NULL`. If this is the case, the pointer has not been initialized.

## Standards and compliance

ANSI SQL – Compliance level: The `text`, `image`, and `unitext` datatypes are Transact-SQL extensions.

## User-defined datatypes

User-defined datatypes are built from the system datatypes and from the `sysname` or `longsysname` user-defined datatypes. After you create a user-defined datatype, you can use it to define columns, parameters, and variables. Objects that are created from user-defined datatypes inherit the rules, defaults, null type, and `IDENTITY` property of the user-defined datatype, as well as inheriting the defaults and null type of the system datatypes on which the user-defined datatype is based.

A user-defined datatype must be created in each database in which it will be used. Create frequently used types in the model database. These types are automatically added to each new database (including tempdb, which is used for temporary tables) as it is created.

Adaptive Server allows you to create user-defined datatypes, based on any system datatype, using `sp_addtype`. You cannot create a user-defined datatype based on another user-defined datatype, such as `timestamp` or the `tid` datatype in the `pubs2` database.

The `sysname` and `longsysname` datatypes are exceptions to this rule. Though `sysname` and `longsysname` are user-defined datatypes, you can use them to build user-defined datatypes.

User-defined datatypes are database objects. Their names are case-sensitive and must conform to the rules for identifiers.

You can bind rules to user-defined datatypes with `sp_bindrule` and bind defaults with `sp_bindefault`.

By default, objects built on a user-defined datatype inherit the user-defined datatype's null type or `IDENTITY` property. You can override the null type or `IDENTITY` property in a column definition.

Use `sp_rename` to rename a user-defined datatype.

Use `sp_droptype` to remove a user-defined datatype from a database.

---

**Note** You cannot drop a datatype that is already in use in a table.

---

Use `sp_help` to display information about the properties of a system datatype or a user-defined datatype. You can also use `sp_help` to display the datatype, length, precision, and scale for each column in a table.

## Standards and compliance

ANSI SQL – Compliance level: User-defined datatypes are a Transact-SQL extension.

# Transact-SQL Functions

This chapter describes the Transact-SQL functions. Functions are used to return information from the database. They are allowed in the select list, in the where clause, and anywhere an expression is allowed. They are often used as part of a stored procedure or program.

Topics	Page
Types of functions	43
Aggregate functions	49
Datatype conversion functions	55
Date functions	64
Mathematical functions	65
Security functions	66
String functions	67
System functions	68
Text and image functions	69

## Types of functions

Table 2-1 lists the different types of Transact-SQL functions and describes the type of information each returns.

**Table 2-1: Types of Transact-SQL functions**

Type of function	Description
Aggregate functions	Generate summary values that appear as new columns or as additional rows in the query results.
Datatype conversion functions	Change expressions from one datatype to another and specify new display formats for date and time information.
Date functions	Perform computations on datetime, smalldatetime, date, and time values and their components, date parts.
Mathematical functions	Commonly needed for operations on mathematical data.
Security functions	Security-related information.
String functions	Operate on binary data, character strings, and expressions.

Type of function	Description
System functions	Retrieves special information from the database and database objects.
Text and image functions	Supply values commonly needed for operations on text, unitext, and image data.

Table 2-2 lists the functions in alphabetical order.

**Table 2-2: List of Transact-SQL functions**

Function	Type	Return value
abs on page 70	Mathematical	The absolute value of an expression.
acos on page 71	Mathematical	The angle (in radians) with a specified cosine.
ascii on page 72	String	The ASCII code for the first character in an expression.
asin on page 73	Mathematical	The angle (in radians) with a specified sine.
atan on page 74	Mathematical	The angle (in radians) with a specified tangent.
atn2 on page 75	Mathematical	The angle (in radians) with specified sine and cosine.
audit_event_name on page 78	Security	A description of an audit event
avg on page 76	Aggregate	The numeric average of all (distinct) values.
biginttohex on page 80	Datatype conversion	Returns the platform-independent hexadecimal equivalent of the specified integer.
case on page 81		Allows SQL expressions to be written for conditional values. case expressions can be used anywhere a value expression can be used.
cast on page 84	Datatype conversion	A specified value, converted to another datatype
ceiling on page 87	Mathematical	The smallest integer greater than or equal to the specified value.
char on page 89	String	The character equivalent of an integer.
charindex on page 93	String	Returns an integer representing the starting position of an expression.
char_length on page 91	String	The number of characters in an expression.
col_length on page 96	System	The defined length of a column.
col_name on page 97	System	The name of the column with specified table and column IDs.
compare on page 98	System	Returns the following values, based on the collation rules that you chose: <ul style="list-style-type: none"> <li>• 1 – indicates that <i>char_expression1</i> is greater than <i>char_expression2</i></li> <li>• 0 – indicates that <i>char_expression1</i> is equal to <i>char_expression2</i></li> <li>• -1 – indicates that <i>char_expression1</i> is less than <i>char_expression2</i></li> </ul>
convert on page 103	Datatype conversion	The specified value, converted to another datatype or a different datetime display format.
cos on page 109	Mathematical	The cosine of the specified angle (in radians).
cot on page 110	Mathematical	The cotangent of the specified angle (in radians).
count on page 111	Aggregate	The number of (distinct) non-null values as an integer.



Function	Type	Return value
count_big on page 113	Aggregate	The number of (distinct) non-null values as a bigint.
current_date on page 115	Date	Returns the current date.
current_time on page 116	Date	Returns the current time.
curunreservedpgs on page 117	System	The number of free pages in the specified disk piece.
data_pages on page 119	System	The number of pages used by the specified table or index.
datalength on page 123	System	The actual length, in bytes, of the specified column or string.
dateadd on page 124	Date	The date produced by adding a given number of years, quarters, hours, or other date parts to the specified date.
datediff on page 127	Date	The difference between two date expressions.
datetime on page 130	Date	The name of the specified part of a date expression.
datepart on page 132	Date	The integer value of the specified part of a date expression.
day on page 136	Date	Returns an integer that represents the day in the datepart of a specified date.
db_id on page 137	System	The ID number of the specified database.
db_name on page 138	System	The name of the database with a specified ID number.
degrees on page 139	Mathematical	The size, in degrees, of an angle with a specified number of radians.
derived_stat on page 140	System	Returns derived statistics for the specified object and index.
difference on page 143	String	The difference between two <code>soundex</code> values.
exp on page 144	Mathematical	The value that results from raising the constant <code>e</code> to the specified power.
floor on page 145	Mathematical	The largest integer that is less than or equal to the specified value.
get_appcontext on page 147	Security	Returns the value of the attribute in a specified context.
getdate on page 149	Date	The current system date and time.
hextobigint on page 153	Datatype conversion	The bigint value equivalent of a hexadecimal string
hextoint on page 154	Datatype conversion	The platform-independent integer equivalent of the specified hexadecimal string.
host_id on page 155	System	Returns the client computer's operating system process ID for the current Adaptive Server client.
host_name on page 156	System	The current host computer name of the client process.
identity_burn_max on page 157		The <code>identity_burn_max</code> value.

Function	Type	Return value
index_col on page 158	System	The name of the indexed column in the specified table or view.
index_colorder on page 159	System	Returns the column order
inttohex on page 160	Datatype conversion	The platform-independent, hexadecimal equivalent of the specified integer.
is_quiesced on page 161		Indicates whether a database is in quiesce database mode. is_quiesced returns 1 if the database is quiesced and 0 if it is not.
is_sec_service_on on page 163	Security	1 if the security service is active; 0 if it is not.
isnull on page 164	System	Substitutes the value specified in <i>expression2</i> when <i>expression1</i> evaluates to NULL.
lct_admin on page 165	System	Manages the last-chance threshold.
left on page 168	String	Returns a specified number of characters on the left end of a character string.
len on page 170	String	Returns the number of characters, not the number of bytes, of a specified string expression, excluding trailing blanks.
license_enabled on page 171	System	1” if the feature’s license is enabled; 0 if it is not.
list_appcontext on page 172	Security	Lists all the attributes of all the contexts in the current session.
lockscheme on page 173	Mathematical	Returns the locking scheme of the specified object as a string.
log on page 174	Mathematical	The natural logarithm of the specified number.
log10 on page 175	Mathematical	The base 10 logarithm of the specified number.
lower on page 176	String	The lowercase equivalent of the specified expression.
ltrim on page 177	String	The specified expression, trimmed of leading blanks
max on page 178	Aggregate	The highest value in a column.
min on page 180	Aggregate	The lowest value in a column.
month on page 181	Date	An integer that represents the month in the datepart of a specified date
mut_excl_roles on page 182	Security	The mutual exclusivity between two roles.
newid on page 183	System	Generates human-readable, globally unique IDs (GUIDs) in two different formats, based on arguments you provide.
next_identity on page 185	System	Retrieves the next identity value that is available for the next insert.
nullif on page 186		Allows SQL expressions to be written for conditional values. nullif expressions can be used anywhere a value expression can be used; alternative for a case expression.
object_id on page 188	System	The object ID of the specified object.

Function	Type	Return value
object_name on page 189	System	The name of the object with the specified object ID.
pagesize on page 190	Mathematical	Returns the page size, in bytes, for the specified object.
partition_id on page 192		Returns the partition ID of the specified data or index partition name.
partition_name on page 193		The explicit name of a new partition, partition_name returns the partition name of the specified data or index partition id.
patindex on page 194	String, Text, Unitext, and Image	The starting position of the first occurrence of a specified pattern.
pi on page 197	Mathematical	The constant value 3.1415926535897936.
power on page 198	Mathematical	The value that results from raising the specified number to a given power.
proc_role on page 199	Security	1 if the user has the correct role to execute the procedure; 0 if the user does not have this role.
radians on page 201	Mathematical	The size, in radians, of an angle with a specified number of degrees.
rand on page 202	Mathematical	A random value between 0 and 1, generated using the specified seed value.
replicate on page 203	String	A string consisting of the specified expression repeated a given number of times.
reserved_pages on page 204	System	The number of pages allocated to the specified table or index.
reverse on page 206	String	The specified string, with characters listed in reverse order.
right on page 207	String	The part of the character expression, starting the specified number of characters from the right.
rm_appcontext on page 209	Security	Removes a specific application context, or all application contexts.
role_contain on page 210	Security	1 if <i>role2</i> contains <i>role1</i> .
role_id on page 211	Security	The system role ID of the role with the name you specify.
role_name on page 212	Security	The name of a role with the system role ID you specify.
round on page 213	Mathematical	The value of the specified number, rounded to a given number of decimal places.
row_count on page 215	System	An estimate of the number of rows in the specified table.
rtrim on page 216	String	The specified expression, trimmed of trailing blanks.
set_appcontext on page 217	Security	Sets an application context name, attribute name, and attribute value for a user session, defined by the attributes of a specified application.
show_role on page 219	Security	The login's currently active roles.

Function	Type	Return value
show_sec_services on page 220	Security	A list of the user's currently active security services.
sign on page 221	Mathematical	The sign (+1 for positive, 0, or -1 for negative) of the specified value.
sin on page 222	Mathematical	The sine of the specified angle (in radians).
sortkey on page 223	System	Values that can be used to order results based on collation behavior, which allows you to work with character collation behaviors beyond the default set of Latin-character-based dictionary sort orders and case or accent sensitivity.
soundex on page 228	String	A 4-character code representing the way an expression sounds.
space on page 229	String	A string consisting of the specified number of single-byte spaces.
square on page 230	Mathematical	Returns the square of a specified value expressed as a float.
sqrt on page 231	Mathematical	The square root of the specified number.
str on page 232	String	The character equivalent of the specified number.
str_replace on page 234	String	Replaces any instances of the second string expression that occur within the first string expression with a third expression.
stuff on page 236	String	The string formed by deleting a specified number of characters from one string and replacing them with another string.
substring on page 238	String	The string formed by extracting a specified number of characters from another string.
sum on page 240	Aggregate	The total of the values.
suser_id	System	The server user's ID number from the syslogins system table.
suser_name on page 243	System	The name of the current server user, or the user where the server user ID is specified.
syb_quit on page 244	System	Terminates the connection.
syb_sendmsg on page 245	System	Sends a message to a User Datagram Protocol (UDP) port.
tan on page 246	Mathematical	The tangent of the specified angle (in radians).
tempdb_id on page 247	System	The database ID of the temporary database assigned to the specified spid
textptr on page 248	Text, Unitext, and Image	The pointer to the first page of the specified text column.
textvalid on page 249	Text and Image	1 if the pointer to the specified text column is valid; 0 if it is not.
to_unichar on page 250	String	A unichar expression having the value of the integer expression.
tran_dumpable_status on page 251	System	Returns a true/false indication of whether dump transaction is allowed.
tsequal on page 252	System	Compares timestamp values to prevent update on a row that has been modified since it was selected for browsing.

<b>Function</b>	<b>Type</b>	<b>Return value</b>
uhighsurr on page 254	String	1 if the Unicode value at position start is the high half of a surrogate pair (which should appear first in the pair); otherwise 0.
ulowsurr on page 255	String	1 if the Unicode value at position start is the low half of a surrogate pair (which should appear second in the pair); otherwise 0.
upper on page 256	String	The uppercase equivalent of the specified string.
uscalar on page 257	String	The Unicode scalar value for the first Unicode character in an expression.
used_pages on page 258	System	The number of pages used by the specified table and its clustered index.
user on page 260	System	The name of the current server user.
user_id on page 261	System	The ID number of the specified user or the current user.
user_name on page 262	System	The name within the database of the specified user or the current user.
valid_name on page 263	System	0 if the specified string is not a valid identifier; a number other than 0 if the string is valid.
valid_user on page 264	System	1 if the specified ID is a valid user or alias in at least one database on this Adaptive Server.
year on page 265	Date	An integer that represents the year in the datepart of a specified date.

The following sections describe the types of functions in detail. The remainder of the chapter contains descriptions of the individual functions in alphabetical order.

## Aggregate functions

The aggregate functions generate summary values that appear as new columns in the query results. The aggregate functions are:

- avg
- count
- count\_big
- max
- min
- sum

Aggregate functions can be used in the select list or the having clause of a select statement or subquery. They cannot be used in a where clause.

Each aggregate in a query requires its own worktable. Therefore, a query using aggregates cannot exceed the maximum number of worktables allowed in a query (12).

When an aggregate function is applied to a char datatype value, it implicitly converts the value to varchar, stripping all trailing blanks. Likewise, a unichar datatype value is implicitly converted to univarchar.

The max, min, and count aggregate functions have semantics that include the unichar datatype.

### Aggregates used with *group by*

Aggregates are often used with group by. With group by, the table is divided into groups. Aggregates produce a single value for each group. Without group by, an aggregate function in the select list produces a single value as a result, whether it is operating on all the rows in a table or on a subset of rows defined by a where clause.

### Aggregate functions and NULL values

Aggregate functions calculate the summary values of the non-null values in a particular column. If the ansinull option is set off (the default), there is no warning when an aggregate function encounters a null. If ansinull is set on, a query returns the following SQLSTATE warning when an aggregate function encounters a null:

```
Warning- null value eliminated in set function
```

### Vector and scalar aggregates

Aggregate functions can be applied to all the rows in a table, in which case they produce a single value, a scalar aggregate. They can also be applied to all the rows that have the same value in a specified column or expression (using the group by and, optionally, the having clause), in which case, they produce a value for each group, a vector aggregate. The results of the aggregate functions are shown as new columns.

You can nest a vector aggregate inside a scalar aggregate. For example:

```
select type, avg(price), avg(avg(price))
from titles
group by type
type
-----
```

UNDECIDED	NULL	15.23
business	13.73	15.23
mod_cook	11.49	15.23
popular_comp	21.48	15.23
psychology	13.50	15.23
trad_cook	15.96	15.23

(6 rows affected)

The group by clause applies to the vector aggregate—in this case, `avg(price)`. The scalar aggregate, `avg(avg(price))`, is the average of the average prices by type in the titles table.

In standard SQL, when a *select\_list* includes an aggregate, all the *select\_list* columns must either have aggregate functions applied to them or be in the group by list. Transact-SQL has no such restrictions.

Example 1 shows a select statement with the standard restrictions. Example 2 shows the same statement with another item (`title_id`) added to the select list. `order by` is also added to illustrate the difference in displays. These “extra” columns can also be referenced in a having clause.

#### Example 1

```
select type, avg(price), avg(advance)
from titles
group by type
type
-----
```

UNDECIDED	NULL	NULL
business	13.73	6,281.25
mod_cook	11.49	7,500.00
popular_comp	21.48	7,500.00
psychology	13.50	4,255.00
trad_cook	15.96	6,333.33

(6 rows affected)

#### Example 2

You can use either a column name or any other expression (except a column heading or alias) after group by.

Null values in the group by column are placed into a single group.

```
select type, title_id, avg(price), avg(advance)
from titles
group by type
order by type
```

type	title_id	avg(price)	avg(advance)
UNDECIDED	MC3026	NULL	NULL
business	BU1032	13.73	6,281.25
business	BU1111	13.73	6,281.25
business	BU2075	13.73	6,281.25
business	BU7832	13.73	6,281.25
mod_cook	MC2222	11.49	7,500.00
mod_cook	MC3021	11.49	7,500.00
popular_comp	PC1035	21.48	7,500.00
popular_comp	PC8888	21.48	7,500.00
popular_comp	PC9999	21.48	7,500.00
psychology	PS1372	13.50	4,255.00
psychology	PS2091	13.50	4,255.00
psychology	PS2106	13.50	4,255.00
psychology	PS3333	13.50	4,255.00
psychology	PS7777	13.50	4,255.00
trad_cook	TC3218	15.96	6,333.33
trad_cook	TC4203	15.96	6,333.33
trad_cook	TC7777	15.96	6,333.33

**Example 3**

The compute clause in a select statement uses row aggregates to produce summary values. The row aggregates make it possible to retrieve detail and summary rows with one command. Example 3 illustrates this feature:

```
select type, title_id, price, advance
from titles
where type = "psychology"
order by type
compute sum(price), sum(advance) by type
```

type	title_id	price	advance
psychology	PS1372	21.59	7,000.00
psychology	PS2091	10.95	2,275.00
psychology	PS2106	7.00	6,000.00
psychology	PS3333	19.99	2,000.00
psychology	PS7777	7.99	4,000.00
		sum	sum
		67.52	21,275.00



Note the difference in display between Example 3 and the examples without compute (Example 1 and Example 2).

You cannot use aggregate functions on virtual tables such as `sysprocesses` and `syslocks`.

If you include an aggregate function in the `select` clause of a cursor, that cursor cannot be updated.

## Aggregate functions as row aggregates

Row aggregate functions generate summary values that appear as additional rows in the query results.

To use the aggregate functions as row aggregates, use the following syntax:

*Start of select statement*

```
compute row_aggregate(column_name)
      [, row_aggregate(column_name)]...
      [by column_name [, column_name]...]
```

Where:

- `column_name` is the name of a column. It must be enclosed in parentheses. Only exact numeric, approximate numeric, and money columns can be used with `sum` and `avg`.

One `compute` clause can apply the same function to several columns. When using more than one function, use more than one `compute` clause.

- `by` indicates that row aggregate values are to be calculated for subgroups. Whenever the value of the `by` item changes, row aggregate values are generated. If you use `by`, you must use `order by`.

Listing more than one item after `by` breaks a group into subgroups and applies a function at each level of grouping.

The row aggregates make it possible to retrieve detail and summary rows with one command. The aggregate functions, on the other hand, ordinarily produce a single value for all the selected rows in the table or for each group, and these summary values are shown as new columns.

The following examples illustrate the differences:

```
select type, sum(price), sum(advance)
from titles
where type like "%cook"
group by type
```

```

type
-----
mod_cook          22.98          15,000.00
trad_cook         47.89          19,000.00

```

(2 rows affected)

```

select type, price, advance
from titles
where type like "%cook"
order by type
compute sum(price), sum(advance) by type

```

```

type          price          advance
-----
mod_cook       2.99          15,000.00
mod_cook      19.99           0.00
              sum              sum
-----
              22.98          15,000.00

```

```

type          price          advance
-----
trad_cook     11.95           4,000.00
trad_cook     14.99           8,000.00
trad_cook     20.95           7,000.00
              sum              sum
-----
              47.89          19,000.00

```

(7 rows affected)

```

type          price          advance
-----
mod_cook       2.99          15,000.00
mod_cook      19.99           0.00

```

Compute Result:

```

-----
              22.98          15,000.00
type          price          advance
-----
trad_cook     11.95           4,000.00
trad_cook     14.99           8,000.00
trad_cook     20.95           7,000.00

```

Compute Result:

```

-----
              47.89          19,000.00
(7 rows affected)

```

The columns in the compute clause must appear in the select list.

The order of columns in the select list overrides the order of the aggregates in the compute clause. For example:

```
create table t1 (a int, b int, c int null)
insert t1 values(1,5,8)
insert t1 values(2,6,9)
```

(1 row affected)

```
compute sum(c), max(b), min(a)
select a, b, c from t1
```

a	b	c
1	5	8
2	6	9

Compute Result:

1	6	17
---	---	----

If the `ansinull` option is set off (the default), there is no warning when a row aggregate encounters a null. If `ansinull` is set on, a query returns the following `SQLSTATE` warning when a row aggregate encounters a null:

```
Warning - null value eliminated in set function
```

You cannot use `select into` in the same statement as a compute clause because there is no way to store the compute clause output in the resulting table.

## Datatype conversion functions

Datatype conversion functions change expressions from one datatype to another and specify new display formats for date and time information. The datatype conversion functions are:

- `cast`
- `convert`
- `inttohex`
- `hextoint`
- `hextobigint`

- biginttohex
- str

You can use the datatype conversion functions in the select list, in the where clause, and anywhere else an expression is allowed.

Adaptive Server performs certain datatype conversions automatically. These are called **implicit conversions**. For example, if you compare a char expression and a datetime expression, or a smallint expression and an int expression, or char expressions of different lengths, Adaptive Server automatically converts one datatype to another.

You must request other datatype conversions explicitly, using one of the built-in datatype conversion functions. For example, before concatenating numeric expressions, you must convert them to character expressions.

Adaptive Server does not allow you to convert certain datatypes to certain other datatypes, either implicitly or explicitly. For example, you cannot convert the following:

- smallint data to datetime
- datetime data to smallint
- binary or varbinary data to smalldatetime or datetime data

Unsupported conversions result in error messages.

Table 2-3 indicates whether individual datatype conversions are performed implicitly, explicitly, or are not supported.

**Table 2-3: Explicit, implicit, and unsupported datatype conversions**

From	binary	varbinary	bit	[n]char	[n]varchar	datetime	smalldatetime	tinyint	smallint	unsigned smallint	int	unsigned int	bigint	unsigned bigint	decimal	numeric	float	real	money	smallmoney	text	unitext	image	unichar	univarchar	date	time
binary	-	I	I	I	I	U	U	I	I	I	I	I	I	I	I	I	I	I	I	I	U	I	I	I	I	I	I
varbinary	I	-	I	I	I	U	U	I	I	I	I	I	I	I	I	I	I	I	I	I	U	I	I	I	I	I	I
bit	I	I	-	I	I	U	U	I	I	I	I	I	I	I	I	I	I	I	I	I	U	U	U	E	E	U	U
[n]char	I	I	E	-	I	I	I	E	E	E	E	E	E	E	E	E	E	E	E	E	I	I	I	I	I	I	I
[n]varchar	I	I	E	I	-	I	I	E	E	E	E	E	E	E	E	E	E	E	E	E	I	I	I	I	I	I	I
datetime	I	I	U	I	I	-	I	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	I	I	I	I
smalldatetime	I	I	U	I	I	I	-	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	I	I	I	I
tinyint	I	I	I	E	E	U	U	-	I	I	I	I	I	I	I	I	I	I	I	I	U	U	U	E	E	U	U
smallint	I	I	I	E	E	U	U	I	-	I	I	I	I	I	I	I	I	I	I	I	U	U	U	U	E	U	U
unsigned smallint	I	I	I	E	E	U	U	I	I	-	I	I	I	I	I	I	I	I	I	I	U	U	U	E	E	U	U
int	I	I	I	E	E	U	U	I	I	I	-	I	I	I	I	I	I	I	I	I	U	U	U	E	E	U	U
unsigned int	I	I	I	E	E	U	U	I	I	I	I	-	I	I	I	I	I	I	I	I	U	U	U	E	E	U	U
bigint	I	I	I	E	E	U	U	I	I	I	I	I	-	I	I	I	I	I	I	I	U	U	U	E	E	U	U
unsigned bigint	I	I	I	E	E	U	U	I	I	I	I	I	I	-	I	I	I	I	I	I	U	U	U	E	E	U	U
decimal	I	I	I	E	E	U	U	I	I	I	I	I	I	I	-	I	I	I	I	I	U	U	U	E	E	U	U
numeric	I	I	I	E	E	U	U	I	I	I	I	I	I	I	I	-	I	I	I	I	U	U	U	E	E	U	U
float	I	I	I	E	E	U	U	I	I	I	I	I	I	I	I	I	-	I	I	I	U	U	U	E	E	U	U
real	I	I	I	E	E	U	U	I	I	I	I	I	I	I	I	I	I	-	I	I	U	U	U	E	E	U	U
money	I	I	I	I	I	U	U	I	I	I	I	I	I	I	I	I	I	I	I	-	I	U	U	E	E	U	U
smallmoney	I	I	I	I	I	U	U	I	I	I	I	I	I	I	I	I	I	I	I	I	-	U	U	E	E	U	U
text	U	U	U	E	E	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	-	I	U	E	E	U	U
unitext	E	E	E	E	E	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	I	-	I	U	U	U	U
image	E	E	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	I	-	E	E	U	U
unichar	I	I	E	I	I	I	I	E	E	E	E	E	E	E	E	E	E	E	E	E	I	I	I	-	I	I	I
univarchar	I	I	E	I	I	I	I	E	E	E	E	E	E	E	E	E	E	E	E	E	I	I	I	I	-	I	I
date	I	I	U	I	I	I	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	I	I	-	I
time	I	I	U	I	I	I	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	U	I	I	I	-

- Datatype conversion key
- E – explicit datatype conversion is required.
  - I – conversion can be done either implicitly, or with an explicit datatype conversion function.
  - I/E – Explicit datatype conversion function required when there is loss of precision or scale, and `arithabortnumeric_truncation` is on; implicit conversion allowed otherwise.
  - U – unsupported conversion.
  - – conversion of a datatype to itself. These conversions are allowed, but are meaningless.

## Converting character data to a noncharacter type

You can convert character data to a noncharacter type—such as a money, date/time, exact numeric, or approximate numeric type—if it consists entirely of characters that are valid for the new type. Leading blanks are ignored. However, if a `char` expression that consists of a blank or blanks is converted to a datetime expression, Adaptive Server converts the blanks into the default datetime value of “Jan 1, 1900.”

Syntax errors are generated when the data includes unacceptable characters. Following are some examples of characters that cause syntax errors:

- Commas or decimal points in integer data
- Commas in monetary data
- Letters in exact or approximate numeric data or bit stream data
- Misspelled month names in date and time data

Implicit conversions between `unichar/univarchar` and `datetime/smalldatetime` are supported.

## Converting from one character type to another

When converting from a multibyte character set to a single-byte character set, characters with no single-byte equivalent are converted to question marks.

text and unitext columns can be explicitly converted to char, nchar, varchar, unichar, univarchar, or nvarchar. You are limited to the maximum length of the character datatypes, which is determined by the maximum column size for your server's logical page size. If you do not specify the length, the converted value has a default length of 30 bytes.

## Converting numbers to a character type

Exact and approximate numeric data can be converted to a character type. If the new type is too short to accommodate the entire string, an insufficient space error is generated. For example, the following conversion tries to store a 5-character string in a 1-character type:

```
select convert(char(1), 12.34)
Insufficient result space for explicit conversion
of NUMERIC value '12.34' to a CHAR field.
```

When converting float data to a character type, the new type should be at least 25 characters long.

---

**Note** The str function may be preferable to convert or cast when making conversions, because it provides more control over conversions and avoids errors.

---

## Rounding during conversion to and from money types

The money and smallmoney types store 4 digits to the right of the decimal point, but round up to the nearest hundredth (.01) for display purposes. When data is converted to a money type, it is rounded up to four places.

Data converted from a money type follows the same rounding behavior if possible. If the new type is an exact numeric with less than three decimal places, the data is rounded to the scale of the new type. For example, when \$4.50 is converted to an integer, it yields 5:

```
select convert(int, $4.50)
-----
5
```

Data converted to money or smallmoney is assumed to be in full currency units such as dollars rather than in fractional units such as cents. For example, the integer value of 5 is converted to the money equivalent of 5 dollars, not 5 cents, in the us\_english language.

## **Converting *date* and *time* information**

Data that is recognizable as a date can be converted to datetime, smalldatetime, date, or time. Incorrect month names lead to syntax errors. Dates that fall outside the acceptable range for the datatype lead to arithmetic overflow errors.

When datetime values are converted to smalldatetime, they are rounded to the nearest minute.

When converting date data to a character type, use style numbers 1 through 7 (101 through 107) or 10 through 12 (110 through 112) in Table 2-6 on page 104 to specify the display format. The default value is 100 (mon dd yyyy hh:miAM (or PM)). If date data is converted to a style that contains a time portion, that time portion reflects the default value of zero.

When converting time data to a character type, use style number 8 or 9 (108 or 109) to specify the display format. The default is 100 (mon dd yyyy hh:miAM (or PM)). If time data is converted to a style that contains a date portion, the default date of Jan 1, 1900 is displayed.

## **Converting between numeric types**

You can convert data from one numeric type to another. Errors can occur if the new type is an exact numeric with precision or scale that is not sufficient to hold the data.

For example, if you provide a float or numeric value as an argument to a built-in function that expects an integer, the value of the float or numeric is truncated. However, Adaptive Server does not implicitly convert numerics that have a fractional part but returns a scale error message. For example, Adaptive Server returns error 241 for numerics that have a fractional part and error 257 if other datatypes are passed.

Use the arithabort and arithignore options to determine how Adaptive Server handles errors resulting from numeric conversions.



## Arithmetic overflow and divide-by-zero errors

Divide-by-zero errors occur when Adaptive Server tries to divide a numeric value by zero. Arithmetic overflow errors occur when the new type has too few decimal places to accommodate the results. This happens during:

- Explicit or implicit conversions to exact types with a lower precision or scale
- Explicit or implicit conversions of data that falls outside the acceptable range for a money or date/time type
- Conversions of hexadecimal strings requiring more than 4 bytes of storage using hexint

Both arithmetic overflow and divide-by-zero errors are considered serious, whether they occur during an implicit or explicit conversion. Use the `arithabort arith_overflow` option to determine how Adaptive Server handles these errors. The default setting, `arithabort arith_overflow on`, rolls back the entire transaction in which the error occurs. If the error occurs in a batch that does not contain a transaction, `arithabort arith_overflow on` does not roll back earlier commands in the batch, and Adaptive Server does not execute statements that follow the error-generating statement in the batch. If you set `arithabort arith_overflow off`, Adaptive Server aborts the statement that causes the error, but continues to process other statements in the transaction or batch. You can use the `@@error` global variable to check statement results.

Use the `arithignore arith_overflow` option to determine whether Adaptive Server displays a message after these errors. The default setting, `off`, displays a warning message when a divide-by-zero error or a loss of precision occurs. Setting `arithignore arith_overflow on` suppresses warning messages after these errors. You can omit optional `arith_overflow` keyword without any effect.

## Scale errors

When an explicit conversion results in a loss of scale, the results are truncated without warning. For example, when you explicitly convert a float, numeric, or decimal type to an integer, Adaptive Server assumes you want the result to be an integer and truncates all numbers to the right of the decimal point.

During implicit conversions to numeric or decimal types, loss of scale generates a scale error. Use the `arithabort numeric_truncation` option to determine how serious such an error is considered. The default setting, `arithabort numeric_truncation on`, aborts the statement that causes the error, but continues to process other statements in the transaction or batch. If you set `arithabort numeric_truncation off`, Adaptive Server truncates the query results and continues processing.

---

**Note** For entry level ANSI SQL compliance, set:

- `arithabort arith_overflow off`
  - `arithabort numeric_truncation on`
  - `arithignore off`
- 

## Domain errors

The `convert` function generates a domain error when the function's argument falls outside the range over which the function is defined. This happens rarely.

## Conversions between binary and integer types

The binary and varbinary types store hexadecimal-like data consisting of a "0x" prefix followed by a string of digits and letters.

These strings are interpreted differently by different platforms. For example, the string "0x0000100" represents 65536 on machines that consider byte 0 most significant (little-endian) and 256 on machines that consider byte 0 least significant (big-endian).

Binary types can be converted to integer types either explicitly, using the `convert` function, or implicitly. If the data is too short for the new type, it is stripped of its "0x" prefix and zero-padded. If it is too long, it is truncated.

Both `convert` and the implicit datatype conversions evaluate binary data differently on different platforms. Because of this, results may vary from one platform to another. Use the `hextoint` function for platform-independent conversion of hexadecimal strings to integers, and the `inttohex` function for platform-independent conversion of integers to hexadecimal values. Use the `hextobigint` function for platform-independent conversion of hexadecimal strings to 64-bit integers, and the `biginttohex` function for platform-independent conversion of 64-bit integers to hexadecimal values.

## Converting between binary and numeric or decimal types

In binary and varbinary data strings, the first two digits after “0x” represent the binary type: “00” represents a positive number and “01” represents a negative number. When you convert a binary or varbinary type to numeric or decimal, be sure to specify the “00” or “01” values after the “0x” digit; otherwise, the conversion will fail.

For example, here is how to convert the following binary data to numeric:

```
select convert(numeric
(38, 18), 0x0000000000000000000000000000000000000000000000000000000000000000)
-----
123.456000
```

This example converts the same numeric data back to binary:

```
select convert(binary, convert(numeric(38, 18), 123.456))
-----
0x0000000000000000000000000000000000000000000000000000000000000000
```

## Converting image columns to binary types

You can use the convert function to convert an image column to binary or varbinary. You are limited to the maximum length of the binary datatypes, which is determined by the maximum column size for your server’s logical page size. If you do not specify the length, the converted value has a default length of 30 characters.

## Converting other types to *bit*

Exact and approximate numeric types can be converted to the bit type implicitly. Character types require an explicit convert function.

The expression being converted must consist only of digits, a decimal point, a currency symbol, and a plus or minus sign. The presence of other characters generates syntax errors.

The bit equivalent of 0 is 0. The bit equivalent of any other number is 1.

## Converting NULL value

You can use the convert function to change NULL to NOT NULL and NOT NULL to NULL.

## Date functions

The date functions manipulate values of the datatypes datetime, smalldatetime, date or time.

You can use date functions in the select list or where clause of a query.

Use the datetime datatype only for dates after January 1, 1753. datetime values must be enclosed in single or double quotes. Use date for dates from January, 1 0001 to January 1, 9999. date values must be enclosed in single or double quotes. Use char, nchar, varchar, or nvarchar for earlier dates. Adaptive Server recognizes a wide variety of date formats. See “Datatype conversion functions” on page 55 and “Date and time datatypes” on page 20 for more information.

Adaptive Server automatically converts between character and datetime values when necessary (for example, when you compare a character value to a datetime value).

The date datatype can cover dates from January 1, 0001 to January 1, 9999.

## Date parts

The date parts, the abbreviations recognized by Adaptive Server, and the acceptable values are:

Date part	Abbreviation	Values
year	yy	1753 – 9999 (2079 for smalldatetime)
quarter	qq	1 – 4
month	mm	1 – 12
week	wk	1 – 54
day	dd	1 – 31
dayofyear	dy	1 – 366
weekday	dw	1 – 7 (Sun. – Sat.)
hour	hh	0 – 23

Date part	Abbreviation	Values
minute	mi	0 – 59
second	ss	0 – 59
millisecond	ms	0 – 999

When you enter a year as two digits (yy):

- Numbers less than 50 are interpreted as 20yy. For example, 01 is 2001, 32 is 2032, and 49 is 2049.
- Numbers equal to or greater than 50 are interpreted as 19yy. For example, 50 is 1950, 74 is 1974, and 99 is 1999.

Milliseconds can be preceded either with a colon or a period. If preceded by a colon, the number means thousandths of a second. If preceded by a period, a single digit means tenths of a second, two digits mean hundredths of a second, and three digits mean thousandths of a second. For example, “12:30:20:1” means twenty and one-thousandth of a second past 12:30; “12:30:20.1” means twenty and one-tenth of a second past 12:30. Adaptive Server may round or truncate millisecond values when inserting datetime or time data, as these datatypes have a granularity of 1/300th of a second rather than 1/1000th of a second. You can use the time datatype for time information.

## Mathematical functions

Mathematical functions return values commonly needed for operations on mathematical data. Mathematical function names are not keywords.

Each function also accepts arguments that can be implicitly converted to the specified type. For example, functions that accept approximate numeric types also accept integer types. Adaptive Server automatically converts the argument to the desired type.

The mathematical functions are:

- |           |              |            |         |
|-----------|--------------|------------|---------|
| • abs     | • cos        | • log      | • rand  |
| • acos    | • cot        | • log10    | • round |
| • asin    | • degrees    | • pagesize | • sign  |
| • atan    | • exp        | • pi       | • sin   |
| • atn2    | • floor      | • power    | • sqrt  |
| • ceiling | • lockscheme | • radians  | • tan   |

Error traps are provided to handle domain or range errors of these functions. Users can set the `arithabort` and `arithignore` options to determine how domain errors are handled:

- `arithabort arith_overflow` specifies behavior following a divide-by-zero error or a loss of precision. The default setting, `arithabort arith_overflow on`, rolls back the entire transaction or aborts the batch in which the error occurs. If you set `arithabort arith_overflow off`, Adaptive Server aborts the statement that causes the error, but continues to process other statements in the transaction or batch.
- `arithabort numeric_truncation` specifies behavior following a loss of scale by an exact numeric type during an implicit datatype conversion. (When an explicit conversion results in a loss of scale, the results are truncated without warning.) The default setting, `arithabort numeric_truncation on`, aborts the statement that causes the error, but continues to process other statements in the transaction or batch. If you set `arithabort numeric_truncation off`, Adaptive Server truncates the query results and continues processing.
- By default, the `arithignore arith_overflow` option is turned off, causing Adaptive Server to display a warning message after any query that results in numeric overflow. Set the `arithignore` option on to ignore overflow errors.

## Security functions

Security functions return security-related information.

The security functions are:

- `is_sec_service_on`
- `show_sec_services`
- `get_appcontext`
- `list_appcontext`
- `set_appcontext`
- `rm_appcontext`
- `show_role`
- `proc_role`
- `role_contain`
- `role_id`
- `role_name`

## String functions

String functions operate on binary data, character strings, and expressions. The string functions are:

- `ascii`
- `char`
- `charindex`
- `char_length`
- `difference`
- `lower`
- `ltrim`
- `patindex`
- `replicate`
- `reverse`
- `right`
- `rtrim`
- `soundex`
- `space`
- `str`
- `stuff`
- `substring`
- `to_unichar`
- `uhighsurr`
- `ulowsurr`
- `upper`
- `uscalar`

You can nest string functions and use them in a select list, in a where clause, or anywhere an expression is allowed. When you use constants with a string function, enclose them in single or double quotes. String function names are not keywords.

Each string function also accepts arguments that can be implicitly converted to the specified type. For example, functions that accept approximate numeric expressions also accept integer expressions. Adaptive Server automatically converts the argument to the desired type.

When a string function accepts two character expressions but only one expression is `unichar`, the other expression is “promoted” and internally converted to `unichar`. This follows existing rules for mixed-mode expressions. However, this conversion may cause truncation, since `unichar` data sometimes takes twice the space.

## Limits on string functions

Results of string functions are limited to 16K. This limit is independent of the server’s page size. In Transact-SQL string functions and string variables, literals can be as large as 16K even on a 2K page size.

If `set string_truncation` is on, a user receives an error if an insert or update truncates a character string. However, Adaptive Server does not report an error if a *displayed* string is truncated. For example:

```
select replicate("a", 16383) + replicate("B", 4000)
```

This shows that the total length would be 20383, but the result string is restricted to 16K.

## System functions

System functions return special information from the database. The system functions are:

- col\_length
- col\_name
- curunreservedpgs
- data\_pages
- datalength
- db\_id
- db\_name
- host\_id
- host\_name
- index\_col
- is\_quiesced
- isnull
- object\_id
- object\_name
- reserved\_pages
- row\_count
- show\_role
- suser\_id
- suser\_name
- tempdb\_id
- tran\_dumptable\_status
- tsequal
- used\_pages
- user
- user\_id
- user\_name
- valid\_name
- valid\_user

The system functions can be used in a select list, in a where clause, and anywhere an expression is allowed.

When the argument to a system function is optional, the current database, host computer, server user, or database user is assumed.

## Text, unitext, and image columns

text, unitext, and image columns cannot be used:

- As parameters to stored procedures
- As values passed to stored procedures
- As local variables
- In order by, compute, and group by clauses
- In an index
- In a where clause clause, except with the keyword like
- In joins

In triggers, both the inserted and deleted text values reference the new value; you cannot reference the old value.



## Text and image functions

Text and image functions operate on text, image, and unitext data. The text and image functions are:

- `textptr`
- `textvalid`

Text and image built-in function names are not keywords. Use the set `textsize` option to limit the amount of text, image, and unitext data that is retrieved by a `select` statement.

You can use the `patindex` text function on text, image, and unitext columns and can consider it on a text and image function.

You can use the `datalength` function to display the length of data in text, image, and unitext columns.

## abs

Description	Returns the absolute value of an expression.
Syntax	<code>abs(<i>numeric_expression</i>)</code>
Parameters	<i>numeric_expression</i> is a column, variable, or expression with datatype that is an exact numeric, approximate numeric, money, or any type that can be implicitly converted to one of these types.
Examples	Returns the absolute value of -1: <pre>select abs(-1) ----- 1</pre>
Usage	<code>abs</code> , a mathematical function, returns the absolute value of a given expression. Results are of the same type and have the same precision and scale as the numeric expression.
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>abs</code> .
See also	“Mathematical functions” on page 65 for general information about mathematical functions. <b>Functions</b> ceiling, floor, round, sign

## acos

Description	Returns the angle (in radians) with a specified cosine.
Syntax	<code>acos(<i>cosine</i>)</code>
Parameters	<i>cosine</i> is the cosine of the angle, expressed as a column name, variable, or constant of type float, real, double precision, or any datatype that can be implicitly converted to one of these types.
Examples	Returns the angle where the cosine is 0.52: <pre>select acos(0.52) ----- 1.023945</pre>
Usage	<code>acos</code> , a mathematical function, returns the angle (in radians) where the cosine is the specified value.
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>acos</code> .
See also	“Mathematical functions” on page 65 for general information about mathematical functions. <b>Functions</b> <code>cos</code> , degrees, radians

## ascii

Description	Returns the ASCII code for the first character in an expression.												
Syntax	<code>ascii(char_expr   uchar_expr)</code>												
Parameters	<p><i>char_expr</i> is a character-type column name, variable, or constant expression of char, varchar, nchar, or nvarchar type.</p> <p><i>uchar_expr</i> is a character-type column name, variable, or constant expression of unichar or univarchar type.</p>												
Examples	<pre>select au_lname, ascii(au_lname) from authors where ascii(au_lname) &lt; 70</pre> <table> <thead> <tr> <th>au_lname</th> <th></th> </tr> </thead> <tbody> <tr> <td>Bennet</td> <td>66</td> </tr> <tr> <td>Blotchet-Halls</td> <td>66</td> </tr> <tr> <td>Carson</td> <td>67</td> </tr> <tr> <td>DeFrance</td> <td>68</td> </tr> <tr> <td>Dull</td> <td>68</td> </tr> </tbody> </table> <p>Returns the author's last names and the ASCII codes for the first letters in their last names, if the ASCII code is less than 70.</p>	au_lname		Bennet	66	Blotchet-Halls	66	Carson	67	DeFrance	68	Dull	68
au_lname													
Bennet	66												
Blotchet-Halls	66												
Carson	67												
DeFrance	68												
Dull	68												
Usage	<ul style="list-style-type: none"> <li>• <code>ascii</code>, a string function, returns the ASCII code for the first character in the expression.</li> <li>• When a string function accepts two character expressions but only one expression is <code>unichar</code>, the other expression is “promoted” and internally converted to <code>unichar</code>. This follows existing rules for mixed-mode expressions. However, this conversion may cause truncation, since <code>unichar</code> data sometimes takes twice the space.</li> <li>• If <i>char_expr</i> or <i>uchar_expr</i> is <code>NULL</code>, returns <code>NULL</code>.</li> </ul>												
Standards	ANSI SQL – Compliance level: Transact-SQL extension.												
Permissions	Any user can execute <code>ascii</code> .												
See also	For general information about string functions, see “String functions” on page 67.												
	<b>Functions</b> char, to_unichar												

## asin

Description	Returns the angle (in radians) with a specified sine.
Syntax	<code>asin(<i>sine</i>)</code>
Parameters	<i>sine</i> is the sine of the angle, expressed as a column name, variable, or constant of type float, real, double precision, or any datatype that can be implicitly converted to one of these types.
Examples	<pre>select asin(0.52) -----            0.546851</pre>
Usage	<ul style="list-style-type: none"><li>• <code>asin</code>, a mathematical function, returns the angle (in radians) with a sine of the specified value.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>asin</code> .
See also	“Mathematical functions” on page 65 for general information about mathematical functions.
	<b>Functions</b> degrees, radians, sin

## atan

Description	Returns the angle (in radians) with a specified tangent.
Syntax	<code>atan(<i>tangent</i>)</code>
Parameters	<i>tangent</i> is the tangent of the angle, expressed as a column name, variable, or constant of type float, real, double precision, or any datatype that can be implicitly converted to one of these types.
Examples	<pre>select atan(0.50) -----            0.463648</pre>
Usage	<ul style="list-style-type: none"><li>atan, a mathematical function, returns the angle (in radians) of a tangent with the specified value.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute atan.
See also	“Mathematical functions” on page 65 for general information about mathematical functions. <b>Functions</b> atn2, degrees, radians, tan

## atn2

Description	Returns the angle (in radians) with specified sine and cosine.
Syntax	<code>atn2(<i>sine</i>, <i>cosine</i>)</code>
Parameters	<p><i>sine</i> is the sine of the angle, expressed as a column name, variable, or constant of type float, real, double precision, or any datatype that can be implicitly converted to one of these types.</p> <p><i>cosine</i> is the cosine of the angle, expressed as a column name, variable, or constant of type float, real, double precision, or any datatype that can be implicitly converted to one of these types.</p>
Examples	<pre>select atn2(.50, .48) -----            0.805803</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>atn2</code>, a mathematical function, returns the angle (in radians) whose sine and cosine are specified.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>atn2</code> .
See also	“Mathematical functions” on page 65 for general information about mathematical functions.
	<b>Functions</b> <code>atan</code> , <code>degrees</code> , <code>radians</code> , <code>tan</code>

## avg

Description	Returns the numeric average of all (distinct) values.
Syntax	<code>avg([all   distinct] <i>expression</i>)</code>
Parameters	<p><code>all</code> applies avg to all values. <code>all</code> is the default.</p> <p><code>distinct</code> eliminates duplicate values before avg is applied. <code>distinct</code> is optional.</p> <p><i>expression</i> is a column name, constant, function, any combination of column names, constants, and functions connected by arithmetic or bitwise operators, or a subquery. With aggregates, an expression is usually a column name. For more information, see “Expressions” on page 275.</p>

**Examples** **Example 1** Calculates the average advance and the sum of total sales for all business books. Each of these aggregate functions produces a single summary value for all of the retrieved rows:

```
select avg(advance), sum(total_sales)
from titles
where type = "business"

-----
                6,281.25      30788
```

**Example 2** Used with a group by clause, the aggregate functions produce single values for each group, rather than for the entire table. This statement produces summary values for each type of book:

```
select type, avg(advance), sum(total_sales)
from titles
group by type

type
-----
UNDECIDED                NULL          NULL
business                  6,281.25     30788
mod_cook                  7,500.00     24278
popular_comp              7,500.00     12875
psychology                 4,255.00      9939
trad_cook                  6,333.33     19566
```

**Example 3** Groups the titles table by publishers and includes only those groups of publishers who have paid more than \$25,000 in total advances and whose books average more than \$15 in price:



```
select pub_id, sum(advance), avg(price)
from titles
group by pub_id
having sum(advance) > $25000 and avg(price) > $15
```

```
pub_id
-----
0877          41,000.00          15.41
1389          30,000.00          18.98
```

**Usage**

- avg, an aggregate function, finds the average of the values in a column. avg can only be used on numeric (integer, floating point, or money) datatypes. Null values are ignored in calculating averages.
- When you average (signed or unsigned) int, smallint, tinyint data, Adaptive Server returns the result as an int value. When you average (signed or unsigned) bigint data, Adaptive Server returns the result as a bigint value. To avoid overflow errors in DB-Library programs, declare variables used for results appropriately.
- You cannot use avg() with the binary datatypes.
- Since the average value is only defined on numeric datatypes, using avg() Unicode expressions generates an error.

**Standards**

ANSI SQL – Compliance level: Transact-SQL extension.

**Permissions**

Any user can execute avg.

**See also**

For general information about aggregate functions, see “Aggregate functions” on page 49.

**Functions** max, min

## audit\_event\_name

Description Returns a description of an audit event.

Syntax audit\_event\_name(event\_id)

Parameters *event\_id*  
is the number of an audit event.

Examples **Example 1** Queries the audit trail for table creation events:

```
select * from audit_data where audit_event_name(event) = "Create Table"
```

**Example 2** Obtains current audit event values. See the Usage section below for a complete list of audit values and their descriptions.

```
create table #tmp(event_id int, description varchar(255))
go
declare @a int
select @a=1
while (@a<120)
begin
    insert #tmp values (@a, audit_event_name(@a))
    select @a=@a + 1
end
select * from #tmp
go
```

```
-----
event_id  description
-----
         1  Ad hoc Audit Record
         2  Alter Database
        ...
       104  Create Index
       105  Drop Index
```

Usage The following lists the ID and name of each of the 111 audit events:

1	Ad Hoc Audit record	38	Execution Of Stored Procedure	74	Auditing Disabled
2	Alter Database	39	Execution Of Trigger	75	NULL
3	Alter table	40	Grant Command	76	SSO Changed Password
4	BCP In	41	Insert Table	79	NULL
5	NULL	42	Insert View	80	Role Check Performed
6	Bind Default	43	Load Database	81	DBCC Command
7	Bind Message	44	Load Transaction	82	Config
8	Bind Rule	45	Log In	83	Online Database
9	Create Database	46	Log Out	84	Setuser Command
10	Create Table	47	Revoke Command	85	User-defined Function Command
11	Create Procedure	48	RPC In	86	Built-in Function
12	Create Trigger	49	RPC Out	87	Disk Release
13	Create Rule	50	Server Boot	88	Set SSA Command
14	Create Default	51	Server Shutdown	90	Connect Command
15	Create Message	52	NULL	91	Reference
16	Create View	53	NULL	92	Command Text
17	Access To Database	54	NULL	93	JCS Install Command
18	Delete Table	55	Role Toggling	94	JCS Remove Command
19	Delete View	56	NULL	95	Unlock Admin Account
20	Disk Init	57	NULL	96	Quiesce Database Command
21	Disk Refit	58	NULL	97	Create SQLJ Function
22	Disk Reinit	59	NULL	98	Drop SQLJ Function
23	Disk Mirror	60	NULL	99	SSL Administration
24	Disk Unmirror	61	Access To Audit Table	100	Disk Resize
25	Disk Remirror	62	Select Table	101	Mount Database
26	Drop Database	63	Select View	102	Unmount Database
27	Drop Table	64	Truncate Table	103	Login Command
28	Drop Procedure	65	NULL	104	Create Index
29	Drop Trigger	66	NULL	105	Drop Index
30	Drop Rule	67	Unbind Default	106	NULL
31	Drop Default	68	Unbind Rule	107	NULL
32	Drop Message	69	Unbind Message	108	NULL
33	Drop View	70	Update Table	109	NULL
34	Dump Database	71	Update View	110	Deploy UDWS
35	Dump Transaction	72	NULL	111	Undeploy UDWS
36	Fatal Error	73	Auditing Enabled		
37	Nonfatal Error				

---

**Note** Adaptive Server does not log events if audit\_event\_name returns NULL.

---

Standards	ANSI SQL – compliance level: Transact-SQL extension.
Permissions	Any user can execute audit_event_name.
See also	<b>Commands</b> select, sp_audit

## biginttohex

**Description** Returns the platform-independent 8 byte hexadecimal equivalent of the specified integer expression.

**Syntax** biginttohex (*integer\_expression*)

**Parameters** *integer\_expression*  
is the integer value to be converted to a hexadecimal string.

**Examples** This example converts the big integer -9223372036854775808 to a hexadecimal string.

```
1> select biginttohex(-9223372036854775808)
2> go
-----
8000000000000000
```

- Usage**
- biginttohex, a datatype conversion function, returns the platform-independent hexadecimal equivalent of an integer, without a “0x” prefix.
  - Use the biginttohex function for platform-independent conversions of integers to hexadecimal strings. biginttohex accepts any expression that evaluates to a bigint. It always returns the same hexadecimal equivalent for a given expression, regardless of the platform on which it is executed.

**See also** **Functions** convert, hextobigint, hextoint, inttohex

## case

Description	Supports conditional SQL expressions; can be used anywhere a value expression can be used.
Syntax	<pre> case     when <i>search_condition</i> then <i>expression</i>     [when <i>search_condition</i> then <i>expression</i>]...     [else <i>expression</i>] end  case and values syntax: case <i>expression</i>     when <i>expression</i> then <i>expression</i>     [when <i>expression</i> then <i>expression</i>]...     [else <i>expression</i>] end </pre>
Parameters	<p><b>case</b> begins the case expression.</p> <p><b>when</b> precedes the search condition or the expression to be compared.</p> <p><b><i>search_condition</i></b> is used to set conditions for the results that are selected. Search conditions for case expressions are similar to the search conditions in a where clause. Search conditions are detailed in the <i>Transact-SQL User's Guide</i>.</p> <p><b>then</b> precedes the expression that specifies a result value of case.</p> <p><b><i>expression</i></b> is a column name, a constant, a function, a subquery, or any combination of column names, constants, and functions connected by arithmetic or bitwise operators. For more information about expressions, see “Expressions” on page 275 in.</p>
Examples	<p><b>Example 1</b> Selects all the authors from the authors table and, for certain authors, specifies the city in which they live:</p> <pre> select au_lname, postalcode,        case            when postalcode = "94705"                then "Berkeley Author"            when postalcode = "94609"                then "Oakland Author"            when postalcode = "94612"                then "Oakland Author" </pre>

```
        when postalcode = "97330"
          then "Corvallis Author"
        end
    from authors
```

**Example 2** Returns the first occurrence of a non-NULL value in either the lowqty or highqty column of the discounts table:

```
select stor_id, discount,
       coalesce (lowqty, highqty)
from discounts
```

You can also use the following format to produce the same result, since coalesce is an abbreviated form of a case expression:

```
select stor_id, discount,
       case
         when lowqty is not NULL then lowqty
         else highqty
       end
from discounts
```

**Example 3** Selects the *titles* and *type* from the titles table. If the book type is UNDECIDED, nullif returns a NULL value:

```
select title,
       nullif(type, "UNDECIDED")
from titles
```

You can also use the following format to produce the same result, since nullif is an abbreviated form of a case expression:

```
select title,
       case
         when type = "UNDECIDED" then NULL
         else type
       end
from titles
```

**Example 4** Produces an error message, because at least one expression must be something other than the null keyword:

```
select price, coalesce (NULL, NULL, NULL)
from titles
```

All result expressions in a CASE expression must not be NULL.

**Example 5** Produces an error message, because at least two expressions must follow coalesce:

```
select stor_id, discount, coalesce (highqty)
```

from discounts

A single coalesce element is illegal in a COALESCE expression.

Usage	<ul style="list-style-type: none"><li>• case expression simplifies standard SQL expressions by allowing you to express a search condition using a when...then construct instead of an if statement.</li><li>• case expressions can be used anywhere an expression can be used in SQL.</li><li>• If your query produces a variety of datatypes, the datatype of a case expression result is determined by datatype hierarchy, as described in “Datatypes of mixed-mode expressions” on page 7 in. If you specify two datatypes that Adaptive Server cannot implicitly convert (for example, char and int), the query fails.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	case permission defaults to all users. No permission is required to use it.
See also	<b>Commands</b> coalesce, nullif, if...else, select, where clause

## cast

Description	Returns the specified value, converted to another datatype. <code>cast</code> can change the nullability of the source expression, and uses the default format for date and time datatypes.
Syntax	<code>cast (expression as datatype [(length   precision[, scale]])</code>
Parameters	<p><i>expression</i></p> <p>is the value to be converted from one datatype or date format to another. It includes columns, constants, functions, any combination of constants, and functions that are connected by arithmetic or bitwise operators or subqueries.</p> <p>When Java is enabled in the database, <i>expression</i> can be a value to be converted to a Java-SQL class.</p> <p>When <code>unichar</code> is used as the destination datatype, the default length of 30 Unicode values is used if no length is specified.</p> <p><i>length</i></p> <p>is an optional parameter used with <code>char</code>, <code>nchar</code>, <code>unichar</code>, <code>univarchar</code>, <code>varchar</code>, <code>nvarchar</code>, <code>binary</code> and <code>varbinary</code> datatypes. If you do not supply a length, Adaptive Server truncates the data to 30 characters for character types and 30 bytes for binary types. The maximum allowable length for character and binary expression is 64K.</p> <p><i>precision</i></p> <p>is the number of significant digits in a numeric or decimal datatype. For float datatypes, precision is the number of significant binary digits in the mantissa. If you do not supply a precision, Adaptive Server uses the default precision of 18 for numeric and decimal datatypes.</p> <p><i>scale</i></p> <p>is the number of digits to the right of the decimal point in a numeric, or decimal datatype. If you do not supply a scale, Adaptive Server uses the default scale of 0.</p>
Examples	<p><b>Example 1</b> Converts the date into a more readable datetime format:</p> <pre>select cast("01/03/63" as datetime) go</pre> <pre>-----           Jan  3 1963 12:00AM</pre> <p>(1 row affected)</p> <p><b>Example 2</b> Converts the <code>total_sales</code> column in the <code>title</code> database to a 12-character column:</p>



## Usage

```
select title, cast(total_sales as char(12))
```

- For more information about datatype conversion, see “Datatype conversion functions” on page 55.
- `cast` generates a domain error when the argument falls outside the range over which the function is defined. This should happen rarely.
- Use `null` or `not null` to specify the nullability of a target column. You can use `null` or `not null` with `select into` to create a new table and change the datatype and nullability of existing columns in the source table.
- You can use `cast` to convert an image column to binary or varbinary. You are limited to the maximum length of the binary datatypes that is determined by the maximum column size for your server’s logical page size. If you do not specify the length, the converted value has a default length of 30 characters.
- You can use `unichar` expressions as a destination datatype, or they can be converted to another datatype. `unichar` expressions can be converted either explicitly between any other datatype supported by the server, or implicitly.
- If you do not specify length when `unichar` is used as a destination type, the default length of 30 Unicode values is used. If the length of the destination type is not large enough to accommodate the given expression, an error message appears.

## Implicit conversion

Implicit conversion between types when the primary fields do not match may cause data truncation, the insertion of a default value, or an error message to be raised. For example, when a datetime value is converted to a date value, the time portion is truncated, leaving only the date portion. If a time value is converted to a datetime value, a default date portion of Jan 1, 1900 is added to the new datetime value. If a date value is converted to a datetime value, a default time portion of 00:00:00:000 is added to the datetime value.

```
DATE -> VARCHAR, CHAR, BINARY, VARBINARY, DATETIME, SMALLDATETIME
TIME -> VARCHAR, CHAR, BINARY, VARBINARY, DATETIME, SMALLDATETIME
VARCHAR, CHAR, BINARY, VARBINARY, DATETIME, SMALLDATETIME -> DATE
VARCHAR, CHAR, BINARY, VARBINARY, DATETIME, SMALLDATETIME -> TIME
```

## Explicit conversion

If you attempt to explicitly convert a date to a datetime, and the value is outside the datetime range such as “Jan 1, 1000” the conversion is not allowed and an informative error message is raised.

```
DATE -> UNICHAR, UNIVARCHAR
```

```
TIME -> UNICHAR, UNIVARCHAR
UNICHAR, UNIVARCHAR -> DATE
UNICHAR, UNIVARCHAR -> TIME
```

#### Conversions involving Java classes

- When Java is enabled in the database, you can use cast to change datatypes in these ways:
  - Convert Java object types to SQL datatypes.
  - Convert SQL datatypes to Java types.
  - Convert any Java-SQL class installed in Adaptive Server to any other Java-SQL class installed in Adaptive Server if the compile-time datatype of the expression (the source class) is a subclass or superclass of the target class.

The result of the conversion is associated with the current database.

Standards

ANSI SQL – Compliance level: ANSI compliant.

Permissions

Any user can execute cast.

## ceiling

Description	Returns the smallest integer greater than or equal to the specified value.
Syntax	<code>ceiling(<i>value</i>)</code>
Parameters	<p><i>value</i></p> <p>is a column, variable, or expression with a datatype is exact numeric, approximate numeric, money, or any type that can be implicitly converted to one of these types.</p>
Examples	<p><b>Example 1</b></p> <pre>select ceiling(123.45) 124</pre> <p><b>Example 2</b></p> <pre>select ceiling(-123.45) -123</pre> <p><b>Example 3</b></p> <pre>select ceiling(1.2345E2) 24.000000</pre> <p><b>Example 4</b></p> <pre>select ceiling(-1.2345E2) -123.000000</pre> <p><b>Example 5</b></p> <pre>select ceiling(\$123.45) 124.00</pre> <p><b>Example 6</b></p> <pre>select discount, ceiling(discount) from salesdetail where title_id = "PS3333" discount ----- 45.000000          45.000000 46.700000          47.000000 46.700000          47.000000 50.000000          50.000000</pre>
Usage	<ul style="list-style-type: none"> <li>ceiling, a mathematical function, returns the smallest integer that is greater than or equal to the specified value. The return value has the same datatype as the value supplied.</li> </ul>

For numeric and decimal values, results have the same precision as the value supplied and a scale of zero.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute ceiling.

See also

For general information about mathematical functions, see “Mathematical functions” on page 65.

**Command** set

**Functions** abs, floor, round, sign

## char

Description	Returns the character equivalent of an integer.
Syntax	<code>char(integer_expr)</code>
Parameters	<i>integer_expr</i> is any integer (tinyint, smallint, or int) column name, variable, or constant expression between 0 and 255.
Examples	<p><b>Example 1</b></p> <pre>select char(42) - *</pre> <p><b>Example 2</b></p> <pre>select xxx = char(65) xxx ---</pre>
Usage	<ul style="list-style-type: none"> <li>char, a string function, converts a single-byte integer value to a character value (char is usually used as the inverse of ascii).</li> <li>char returns a char datatype. If the resulting value is the first byte of a multibyte character, the character may be undefined.</li> <li>If <i>char_expr</i> is NULL, returns NULL.</li> </ul> <p>Reformatting output with char</p> <ul style="list-style-type: none"> <li>You can use concatenation and char values to add tabs or carriage returns to reformat output. char(10) converts to a return; char(9) converts to a tab. For example:</li> </ul> <pre>/* just a space */ select title_id + " " + title from titles where title_id = "T67061" /* a return */ select title_id + char(10) + title from titles where title_id = "T67061" /* a tab */ select title_id + char(9) + title from titles where title_id = "T67061" ----- T67061 Programming with Courses ----- T67061  Programming with Courses -----</pre>

T67061      Programming with Curses

Standards      ANSI SQL – Compliance level: Transact-SQL extension.

Permissions      Any user can execute char.

See also      For general information about string functions, see “String functions” on page 67.

**Functions**      ascii, str

## char\_length

Description	Returns the number of characters in an expression.
Syntax	<code>char_length(char_expr   uchar_expr)</code>
Parameters	<p><i>char_expr</i> is a character-type column name, variable, or constant expression of char, varchar, nchar, or nvarchar type.</p> <p><i>uchar_expr</i> is a character-type column name, variable, or constant expression of unichar or univarchar type.</p>

### Examples

#### Example 1

```
select char_length(notes) from titles
where title_id = "PC9999"

-----
                39
```

#### Example 2

```
declare @var1 varchar(20), @var2 varchar(20), @char
char(20)
select @var1 = "abcd", @var2 = "abcd  ",
       @char = "abcd"
select char_length(@var1), char_length(@var2),
       char_length(@char)

-----  -----  -----
                4                8                20
```

### Usage

- `char_length`, a string function, returns an integer representing the number of characters in a character expression or text value.
- For variable-length columns and variables, `char_length` returns the number of characters (not the defined length of the column or variable). If explicit trailing blanks are included in variable-length variables, they are not stripped. For literals and fixed-length character columns and variables, `char_length` does not strip the expression of trailing blanks (see Example 2).
- For `unitext`, `unichar`, and `univarchar` columns, `char_length` returns the number of Unicode values (16-bit), with one surrogate pair counted as two Unicode values. For example, this is what is returned if a `unitext` column `ut` contains row value `U+0041U+0042U+d800dc00`:

```
select char_length(ut) from unitable
-----
```

4

- For multibyte character sets, the number of characters in the expression is usually fewer than the number of bytes; use `datalength` to determine the number of bytes.
- For Unicode expressions, returns the number of Unicode values (not bytes) in an expression. Surrogate pairs count as two Unicode values.
- If *char\_expr* or *uchar\_expr* is NULL, `char_length` returns NULL.
- For general information about string functions, see “String functions” on page 67.

Standards                   ANSI SQL – Compliance level: Transact-SQL extension.

Permissions               Any user can execute `char_length`.

See also                   **Function**   `datalength`



## charindex

Description	Returns an integer representing the starting position of an expression.
Syntax	<code>charindex(<i>expression1</i>, <i>expression2</i>)</code>
Parameters	<p><i>expression</i></p> <p>is a binary or character column name, variable, or constant expression. Can be char, varchar, nchar, nvarchar, unichar or univarchar, binary, or varbinary.</p>
Examples	<p>Returns the position at which the character expression “wonderful” begins in the notes column of the titles table:</p> <pre> select charindex("wonderful", notes) from titles where title_id = "TC3218"  -----                 46 </pre>
Usage	<ul style="list-style-type: none"> <li>• <code>charindex</code>, a string function, searches <i>expression2</i> for the first occurrence of <i>expression1</i> and returns an integer representing its starting position. If <i>expression1</i> is not found, <code>charindex</code> returns 0.</li> <li>• If <i>expression1</i> contains wildcard characters, <code>charindex</code> treats them as literals.</li> <li>• If <i>expression2</i> is NULL, returns 0.</li> <li>• If a varchar expression is given as one parameter and a unichar expression as the other, the varchar expression is implicitly converted to unichar (with possible truncation).</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>charindex</code> .
See also	For general information about string functions, see “String functions” on page 67.
	<b>Function</b> patindex

## coalesce

Description	Supports conditional SQL expressions; can be used anywhere a value expression can be used; alternative for a case expression.
Syntax	<code>coalesce(expression, expression [, expression]...)</code>
Parameters	<code>coalesce</code> evaluates the listed expressions and returns the first non-null value. If all expressions are null, <code>coalesce</code> returns NULL.  <i>expression</i> is a column name, a constant, a function, a subquery, or any combination of column names, constants, and functions connected by arithmetic or bitwise operators. For more information about expressions, see “Expressions” on page 275.
Examples	<p><b>Example 1</b> Returns the first occurrence of a non-null value in either the <code>lowqty</code> or <code>highqty</code> column of the <code>discounts</code> table:</p> <pre>select stor_id, discount,        coalesce (lowqty, highqty) from discounts</pre> <p><b>Example 2</b> An alternative way of writing Example 1:</p> <pre>select stor_id, discount,        case          when lowqty is not NULL then lowqty          else highqty        end from discounts</pre>
Usage	<ul style="list-style-type: none"><li>• <code>coalesce</code> expression simplifies standard SQL expressions by allowing you to express a search condition as a simple comparison instead of using a <code>when...then</code> construct.</li><li>• You can use <code>coalesce</code> expressions anywhere an expression in SQL.</li><li>• At least one result of the <code>coalesce</code> expression must return a non-null value. This example produces the following error message: <pre>select price, coalesce (NULL, NULL, NULL) from titles</pre>All result expressions in a CASE expression must not be NULL.</li></ul>

- If your query produces a variety of datatypes, the datatype of a case expression result is determined by datatype hierarchy, as described in “Datatypes of mixed-mode expressions” on page 7. If you specify two datatypes that Adaptive Server cannot implicitly convert (for example, char and int), the query fails.
- `coalesce` is an abbreviated form of a case expression. Example 2 describes an alternative way of writing the `coalesce` statement.
- `coalesce` must be followed by at least two expressions. This example produces the following error message:

```
select stor_id, discount, coalesce (highqty)
from discounts
```

A single `coalesce` element is illegal in a `COALESCE` expression.

Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>coalesce</code> .
See also	<b>Commands</b> <code>case</code> , <code>nullif</code> , <code>select</code> , <code>if...else</code> , <code>where</code> clause

## col\_length

Description	Returns the defined length of a column.
Syntax	<code>col_length(object_name, column_name)</code>
Parameters	<p><i>object_name</i> is name of a database object, such as a table, view, procedure, trigger, default, or rule. The name can be fully qualified (that is, it can include the database and owner name). It must be enclosed in quotes.</p> <p><i>column_name</i> is the name of the column.</p>
Examples	<p>Finds the length of the title column in the titles table. The “x” gives a column heading to the result:</p> <pre>select x = col_length("titles", "title")       x -----       80</pre>
Usage	<ul style="list-style-type: none"><li>• col_length, a system function, returns the defined length of column.</li><li>• For general information about system functions, see “System functions” on page 68.</li><li>• To find the actual length of the data stored in each row, use datalength.</li><li>• For text, unitext, and image columns, col_length returns 16, the length of the binary(16) pointer to the actual text page.</li><li>• For unichar columns, the defined length is the number of Unicode values declared when the column was defined (not the number of bytes represented).</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute col_length.
See also	<b>Function</b> datalength

## col\_name

Description	Returns the name of the column where the table and column IDs are specified, and can be up to 255 bytes in length.
Syntax	<code>col_name(object_id, column_id [, database_id])</code>
Parameters	<p><i>object_id</i> is a numeric expression that is an object ID for a table, view, or other database object. These are stored in the id column of sysobjects.</p> <p><i>column_id</i> is a numeric expression that is a column ID of a column. These are stored in the colid column of syscolumns.</p> <p><i>database_id</i> is a numeric expression that is the ID for a database. These are stored in the db_id column of sysdatabases.</p>
Examples	<pre>select col_name(208003772, 2) ----- title</pre>
Usage	<ul style="list-style-type: none"> <li>col_name, a system function, returns the column's name.</li> <li>For general information about system functions, see "System functions" on page 68.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute col_name.
See also	<b>Functions</b> db_id, object_id

## compare

Description	Allows you to directly compare two character strings based on alternate collation rules.
Syntax	<code>compare ({<i>char_expression1</i> <i>uchar_expression1</i>},           {<i>char_expression2</i> <i>uchar_expression2</i>}),           [<i>{collation_name   collation_ID}</i>]</code>
Parameters	<p><i>char_expression1</i> or <i>uchar_expression1</i> are the character expressions to compare to <i>char_expression2</i> or <i>uchar_expression2</i>.</p> <p><i>char_expression2</i> or <i>uchar_expression2</i> are the character expressions against which to compare <i>char_expression1</i> or <i>uchar_expression1</i>.</p> <p><i>char_expression1</i> and <i>char_expression2</i> can be:</p> <ul style="list-style-type: none"><li>• Character type (char, varchar, nchar, or nvarchar)</li><li>• Character variable, or</li><li>• Constant character expression, enclosed in single or double quotation marks</li></ul> <p><i>uchar_expression1</i> and <i>uchar_expression2</i> can be:</p> <ul style="list-style-type: none"><li>• Character type (unichar or univarchar)</li><li>• Character variable, or</li><li>• Constant character expression, enclosed in single or double quotation marks</li></ul> <p><i>collation_name</i> can be a quoted string or a character variable that specifies the collation to use. Table 2-5 on page 101 shows the valid values.</p> <p><i>collation_ID</i> is an integer constant or a variable that specifies the collation to use. Table 2-5 on page 101 shows the valid values.</p>
Examples	<p><b>Example 1</b> Compares aaa and bbb:</p> <pre>1&gt; select compare ("aaa", "bbb") 2&gt; go  ----- -1 (1 row affected)</pre>

Alternatively, you can also compare aaa and bbb using this format:

```
1> select compare (("aaa"), ("bbb"))
2> go

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

**Example 2** Compares aaa and bbb and specifies binary sort order:

```
1> select compare ("aaa","bbb","binary")
2> go

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

Alternatively, you can compare aaa and bbb using this format, and the collation ID instead of the collation name:

```
1> select compare (("aaa"), ("bbb"), (50))
2> go

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

#### Usage

- The compare function returns the following values, based on the collation rules that you chose:
  - 1 – indicates that *char\_expression1* or *uchar\_expression1* is greater than *char\_expression2* or *uchar\_expression2*.
  - 0 – indicates that *char\_expression1* or *uchar\_expression1* is equal to *char\_expression2* or *uchar\_expression2*.
  - -1 – indicates that *char\_expression1* or *uchar\_expression1* is less than *char\_expression2* or *uchar\_expression2*.
- compare can generate up to six bytes of collation information for each input character. Therefore, the result from using compare may exceed the length limit of the varbinary datatype. If this happens, the result is truncated to fit. Adaptive Server issues a warning message, but the query or transaction that contained the compare function continues to run. Since this limit is dependent on the logical page size of your server, truncation removes result bytes for each input character until the result string is less than the following for DOL and APL tables:

**Table 2-4: Maximum row and column length—APL and DOL**

Locking scheme	Page size	Maximum row length	Maximum column length
APL tables	2K (2048 bytes)	1962	1960 bytes
	4K (4096 bytes)	4010	4008 bytes
	8K (8192 bytes)	8106	8104 bytes
	16K (16384 bytes)	16298	16296 bytes
DOL tables	2K (2048 bytes)	1964	1958 bytes
	4K (4096 bytes)	4012	4006 bytes
	8K (8192 bytes)	8108	8102 bytes
	16K (16384 bytes)	16300	16294 bytes if table does not include any variable length columns
	16K (16384 bytes)	16300 (subject to a max start offset of varlen = 8191)	8191-6-2 = 8183 bytes if table includes at least on variable length column.*

\* This size includes six bytes for the row overhead and two bytes for the row length field

- Both *char\_expression1*, *uchar\_expression1*, and *char\_expression2*, *uchar\_expression2* must be characters that are encoded in the server's default character set.
- *char\_expression1*, *uchar\_expression 1*, or *char\_expression2*, *uchar\_expression2*, or both, can be empty strings:
  - If *char\_expression2* or *uchar\_expression2* is empty, the function returns 1.
  - If both strings are empty, then they are equal, and the function returns 0.
  - If *char\_expression1* or *uchar\_expression 1* is empty, the function returns -1.

The compare function does not equate empty strings and strings containing only spaces. compare uses the sortkey function to generate collation keys for comparison. Therefore, a truly empty string, a string with one space, or a string with two spaces do not compare equally.

- If either *char\_expression1*, *uchar\_expression1*; or *char\_expression2*, *uchar\_expression2* is NULL, then the result is NULL.
- If a varchar expression is given as one parameter and a unichar expression is given as the other, the varchar expression is implicitly converted to unichar (with possible truncation).
- If you do not specify a value for *collation\_name* or *collation\_ID*, compare assumes binary collation.



- Table 2-5 lists the valid values for *collation\_name* and *collation\_ID*.

**Table 2-5: Collation names and IDs**

Description	Collation name	Collation ID
Default Unicode multilingual	default	20
Thai dictionary order	thaidict	21
ISO14651 standard	iso14651	22
UTF-16 ordering – matches UTF-8 binary ordering	utf8bin	24
CP 850 Alternative – no accent	altnoacc	39
CP 850 Alternative – lowercase first	altdict	45
CP 850 Western European – no case preference	altnocsp	46
CP 850 Scandinavian – dictionary ordering	scandict	47
CP 850 Scandinavian – case-insensitive with preference	scannocp	48
GB Pinyin	gbpinyin	n/a
Binary sort	binary	50
Latin-1 English, French, German dictionary	dict	51
Latin-1 English, French, German no case	nocase	52
Latin-1 English, French, German no case, preference	nocasep	53
Latin-1 English, French, German no accent	noaccent	54
Latin-1 Spanish dictionary	espdict	55
Latin-1 Spanish no case	espnocs	56
Latin-1 Spanish no accent	espnoac	57
ISO 8859-5 Russian dictionary	rusdict	58
ISO 8859-5 Russian no case	rusnocs	59
ISO 8859-5 Cyrillic dictionary	cyrdict	63
ISO 8859-5 Cyrillic no case	cyrnocs	64
ISO 8859-7 Greek dictionary	elldict	65
ISO 8859-2 Hungarian dictionary	hundict	69
ISO 8859-2 Hungarian no accents	hunnoac	70
ISO 8859-2 Hungarian no case	hunnocs	71
ISO 8859-9 Turkish dictionary	turdict	72
ISO 8859-9 Turkish no accents	turknoac	73
ISO 8859-9 Turkish no case	turknocs	74
CP932 binary ordering	cp932bin	129
Chinese phonetic ordering	dynix	130
GB2312 binary ordering	gb2312bn	137
Common Cyrillic dictionary	cyrdict	140
Turkish dictionary	turdict	155

<b>Description</b>	<b>Collation name</b>	<b>Collation ID</b>
EUCKSC binary ordering	euckscbn	161
Chinese phonetic ordering	gbpinyin	163
Russian dictionary ordering	rusdict	165
SJIS binary ordering	sjisbin	179
EUCJIS binary ordering	eucjisbn	192
BIG5 binary ordering	big5bin	194
Shift-JIS binary order	sjisbin	259

Standards                   ANSI SQL – Compliance level: Transact-SQL extension.

Permissions                Any user can execute compare.

See also                    **Function**   sortkey

## convert

Description	Returns the specified value, converted to another datatype or a different datetime display format.
Syntax	convert ( <i>datatype</i> [( <i>length</i> )   ( <i>precision</i> [, <i>scale</i> ]]) [null   not null], <i>expression</i> [, <i>style</i> ])
Parameters	<p><i>datatype</i></p> <p>is the system-supplied datatype (for example, char(10), unichar (10), varbinary (50), or int) into which to convert the expression. You cannot use user-defined datatypes.</p> <p>When Java is enabled in the database, <i>datatype</i> can also be a Java-SQL class in the current database.</p> <p><i>length</i></p> <p>is an optional parameter used with char, nchar, unichar, univarchar, varchar, nvarchar, binary, and varbinary datatypes. If you do not supply a length, Adaptive Server truncates the data to 30 characters for the character types and 30 bytes for the binary types. The maximum allowable length for character and binary expression is 64K.</p> <p><i>precision</i></p> <p>is the number of significant digits in a numeric or decimal datatype. For float datatypes, precision is the number of significant binary digits in the mantissa. If you do not supply a precision, Adaptive Server uses the default precision of 18 for numeric and decimal datatypes.</p> <p><i>scale</i></p> <p>is the number of digits to the right of the decimal point in a numeric, or decimal datatype. If you do not supply a scale, Adaptive Server uses the default scale of 0.</p> <p>null   not null</p> <p>specifies the nullability of the result expression. If you do not supply either null or not null, the converted result has the same nullability as the expression.</p> <p><i>expression</i></p> <p>is the value to be converted from one datatype or date format to another.</p> <p>When Java is enabled in the database, <i>expression</i> can be a value to be converted to a Java-SQL class.</p> <p>When unichar is used as the destination datatype, the default length of 30 Unicode values is used if no length is specified.</p>

*style*

is the display format to use for the converted data. When converting money or *smallmoney* data to a character type, use a *style* of 1 to display a comma after every 3 digits.

When converting *datetime* or  *smalldatetime* data to a character type, use the style numbers in Table 2-6 to specify the display format. Values in the left-most column display 2-digit years (yy). For 4-digit years (yyyy), add 100, or use the value in the middle column.

When converting date data to a character type, use style numbers 1 through 7 (101 through 107) or 10 through 12 (110 through 112) in Table 2-6 to specify the display format. The default value is 100 (mon dd yyyy hh:miAM (or PM)). If date data is converted to a style that contains a time portion, that time portion reflects the default value of zero.

When converting time data to a character type, use style number 8 or 9 (108 or 109) to specify the display format. The default is 100 (mon dd yyyy hh:miAM (or PM)). If time data is converted to a style that contains a date portion, the default date of Jan 1, 1900 is displayed.

**Table 2-6: Date format conversions using the style parameter**

Without century (yy)	With century (yyyy)	Standard	Output
-	0 or 100	Default	<i>mon dd yyyy hh:mm AM (or PM)</i>
1	101	USA	<i>mm/dd/yy</i>
2	2	SQL standard	<i>yy.mm.dd</i>
3	103	English/French	<i>dd/mm/yy</i>
4	104	German	<i>dd.mm.yy</i>
5	105		<i>dd-mm-yy</i>
6	106		<i>dd mon yy</i>
7	107		<i>mon dd, yy</i>
8	108		<i>HH:mm:ss</i>
-	9 or 109	Default + milliseconds	<i>mon dd yyyy hh:mm:sss AM (or PM)</i>
10	110	USA	<i>mm-dd-yy</i>
11	111	Japan	<i>yy/mm/dd</i>
12	112	ISO	<i>yy/mm/dd</i>
13	113		<i>yy/mm/dd</i>
14	114		<i>yy/mm/dd</i>

**Key** “mon” indicates a month spelled out, “mm” the month number or minutes. “HH” indicates a 24-hour clock value, “hh” a 12-hour clock value. The last row, 23, includes a literal “T” to separate the date and time portions of the format.

Without century (yy)	With century (yyyy)	Standard	Output
14	114		hh:mi:ss:mmmAM(or PM)
15	115		dd/yy/mm
16	116		mon dd yy HH:mm:ss
17	117		hh:mmAM
18	118		HH:mm
19	119		hh:mm:ss:zzzAM
20	120		hh:mm:ss:zzz
21	121		yy/mm/dd
22	122		yy/mm/dd
23	123		yyyy-mm-ddTHH:mm:ss

**Key** “mon” indicates a month spelled out, “mm” the month number or minutes. “HH” indicates a 24-hour clock value, “hh” a 12-hour clock value. The last row, 23, includes a literal “T” to separate the date and time portions of the format.

The default values (*style* 0 or 100), and *style* 9 or 109 return the century (yyyy). When converting to char or varchar from smalldatetime, styles that include seconds or milliseconds show zeros in those positions.

## Examples

### Example 1

```
select title, convert(char(12), total_sales)
from titles
```

### Example 2

```
select title, total_sales
from titles
where convert(char(20), total_sales) like "1%"
```

### Example 3

Converts the current date to style 3, dd/mm/yy:

```
select convert(char(12), getdate(), 3)
```

### Example 4

If the value pubdate can be null, you must use varchar rather than char, or errors may result:

```
select convert(varchar(12), pubdate, 3) from titles
```

### Example 5

Returns the integer equivalent of the string “0x00000100”. Results can vary from one platform to another:

```
select convert(integer, 0x00000100)
```

### Example 6

Returns the platform-specific bit pattern as a Sybase binary type:

```
select convert (binary, 10)
```

**Example 7** Returns 1, the bit string equivalent of \$1.11:

```
select convert(bit, $1.11)
```

**Example 8** Creates #tempsales with total\_sales of datatype char(100), and does not allow null values. Even if titles.total\_sales was defined as allowing nulls, #tempsales is created with #tempsales.total\_sales not allowing null values:

```
select title, convert(char(100) not null, total_sales)
into #tempsales
from titles
```

## Usage

- convert, a datatype conversion function, converts between a wide variety of datatypes and reformats date/time and money data for display purposes.
- For more information about datatype conversion, see “Datatype conversion functions” on page 55.
- convert – returns the specified value, converted to another datatype or a different datetime display format. When converting from untext to other character and binary datatypes, the result is limited to the maximum length of the destination datatype. If the length is not specified, the converted value has a default size of 30 bytes. If you are using enabled enable surrogate processing, a surrogate pair is returned as a whole. For example, this is what is returned if you convert a untext column that contains data U+0041U+0042U+20acU+0043 (stands for “AB €”) to a UTF-8 varchar(3) column:

```
select convert(varchar(3), ut) from untable
---
AB
```

- convert() generates a domain error when the argument falls outside the range over which the function is defined. This should happen rarely.
- Use null or not null to specify the nullability of a target column. Specifically, this can be used with select into to create a new table and change the datatype and nullability of existing columns in the source table (See Example 8, above).

The result is an undefined value if:

- The expression being converted is to a not null result.
- The expression’s value is null.

Use the following select statement to generate a known non-NULL value for predictable results:

```
select convert(int not null isnull(col2, 5)) from table1
```

- You can use `convert` to convert an image column to binary or varbinary. You are limited to the maximum length of the binary datatypes, which is determined by the maximum column size for your server's logical page size. If you do not specify the length, the converted value has a default length of 30 characters.
- You can use `unichar` expressions as a destination datatype or you can convert them to another datatype. `unichar` expressions can be converted either explicitly between any other datatype supported by the server, or implicitly.
- If you do not specify the length when `unichar` is used as a destination type, the default length of 30 Unicode values is used. If the length of the destination type is not large enough to accommodate the given expression, an error message appears.

#### Implicit conversion

Implicit conversion between types when the primary fields do not match may cause data truncation, the insertion of a default value, or an error message to be raised. For example, when a `datetime` value is converted to a `date` value, the time portion is truncated, leaving only the date portion. If a time value is converted to a `datetime` value, a default date portion of Jan 1, 1900 is added to the new `datetime` value. If a `date` value is converted to a `datetime` value, a default time portion of 00:00:00:000 is added to the `datetime` value.

```
DATE -> VARCHAR, CHAR, BINARY, VARBINARY, DATETIME, SMALLDATETIME
TIME -> VARCHAR, CHAR, BINARY, VARBINARY, DATETIME, SMALLDATETIME
VARCHAR, CHAR, BINARY, VARBINARY, DATETIME, SMALLDATETIME -> DATE
VARCHAR, CHAR, BINARY, VARBINARY, DATETIME, SMALLDATETIME -> TIME
```

#### Explicit conversion

If you attempt to explicitly convert a `date` to a `datetime` and the value is outside the `datetime` range, such as "Jan 1, 1000" the conversion is not allowed and an informative error message is raised.

```
DATE -> UNICHAR, UNIVARCHAR
TIME -> UNICHAR, UNIVARCHAR
UNICHAR, UNIVARCHAR -> DATE
UNICHAR, UNIVARCHAR -> TIME
```

#### Conversions involving Java classes

- When Java is enabled in the database, you can use `convert` to change datatypes in these ways:
  - Convert Java object types to SQL datatypes.
  - Convert SQL datatypes to Java types.

- Convert any Java-SQL class installed in Adaptive Server to any other Java-SQL class installed in Adaptive Server if the compile-time datatype of the expression (the source class) is a subclass or superclass of the target class.

The result of the conversion is associated with the current database.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute convert.

See also

**Documents** *Java in Adaptive Server Enterprise* for a list of allowed datatype mappings and more information about datatype conversions involving Java classes.

**Datatypes** User-defined datatypes

**Functions** hextoint, inttohex



## COS

Description	Returns the cosine of the specified angle.
Syntax	<code>cos(<i>angle</i>)</code>
Parameters	<i>angle</i> is any approximate numeric (float, real, or double precision) column name, variable, or constant expression.
Examples	<pre>select cos(44) 0.999843</pre>
Usage	<ul style="list-style-type: none"><li>• <code>cos</code>, a mathematical function, returns the cosine of the specified angle, in radians.</li><li>• For general information about mathematical functions, see “Mathematical functions” on page 65.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>cos</code> .
See also	<b>Functions</b> <code>acos</code> , <code>degrees</code> , <code>radians</code> , <code>sin</code>

## cot

Description	Returns the cotangent of the specified angle.
Syntax	<code>cot(<i>angle</i>)</code>
Parameters	<i>angle</i> is any approximate numeric (float, real, or double precision) column name, variable, or constant expression.
Examples	<pre>select cot(90) -----           -0.501203</pre>
Usage	<ul style="list-style-type: none"><li>• cot, a mathematical function, returns the cotangent of the specified angle, in radians.</li><li>• For general information about mathematical functions, see “Mathematical functions” on page 65.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute cot.
See also	<b>Functions</b> degrees, radians, sin

## count

Description	Returns the number of (distinct) non-null values, or the number of selected rows as an integer.
Syntax	<code>count([all   distinct] <i>expression</i>)</code>
Parameters	<p><b>all</b> applies count to all values. all is the default.</p> <p><b>distinct</b> eliminates duplicate values before count is applied. distinct is optional.</p> <p><b><i>expression</i></b> is a column name, constant, function, any combination of column names, constants, and functions connected by arithmetic or bitwise operators, or a subquery. With aggregates, an expression is usually a column name. For more information, see “Expressions” on page 275.</p>
Examples	<p><b>Example 1</b> Finds the number of different cities in which authors live:</p> <pre>select count(distinct city) from authors</pre> <p><b>Example 2</b> Lists the types in the titles table, but eliminates the types that include only one book or none:</p> <pre>select type from titles group by type having count(*) &gt; 1</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>count</code>, an aggregate function, finds the number of non-null values in a column. For general information about aggregate functions, see “Aggregate functions” on page 49.</li> <li>• When <code>distinct</code> is specified, <code>count</code> finds the number of unique non-null values. <code>count</code> can be used with all datatypes, including <code>unichar</code>, but cannot be used with <code>text</code> and <code>image</code>. Null values are ignored when counting.</li> <li>• <code>count(column_name)</code> returns a value of 0 on empty tables, on columns that contain only null values, and on groups that contain only null values.</li> <li>• <code>count(*)</code> finds the number of rows. <code>count(*)</code> does not take any arguments, and cannot be used with <code>distinct</code>. All rows are counted, regardless of the presence of null values.</li> </ul>

- When tables are being joined, include count(\*) in the **select list** to produce the count of the number of rows in the joined results. If the objective is to count the number of rows from one table that match criteria, use count(*column\_name*).

- You can use count as an existence check in a subquery. For example:

```
select * from tab where 0 <
      (select count(*) from tab2 where ...)
```

However, because count() counts all matching values, exists or in may return results faster. For example:

```
select * from tab where exists
      (select * from tab2 where ...)
```

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute count.

See also

**Commands** compute clause, group by and having clauses, select, where clause

## count\_big

Description	Returns the number of (distinct) non-null values or the number of selected rows as a bigint.
Syntax	<code>count_big([all   distinct] <i>expression</i>)</code>
Parameters	<p><b>all</b> applies count_big to all values. all is the default.</p> <p><b>distinct</b> eliminates duplicate values before count_big is applied. distinct is optional.</p> <p><b><i>expression</i></b> is a column name, constant, function, any combination of column names, constants, and functions connected by arithmetic or bitwise operators, or a subquery. With aggregates, an expression is usually a column name.</p>
Examples	<p>Finds the number of occurrences of <i>name</i> in systypes:</p> <pre>1&gt; select count_big(name) from systypes 2&gt; go ----- 42</pre>
Usage	<ul style="list-style-type: none"> <li>count_big, an aggregate function, finds the number of non-null values in a column.</li> <li>When distinct is specified, count_big finds the number of unique non-null values. Null values are ignored when counting.</li> <li>count_big(<i>column_name</i>) returns a value of 0 on empty tables, on columns that contain only null values, and on groups that contain only null values.</li> <li>count_big(*) finds the number of rows. count_big(*) does not take any arguments, and cannot be used with distinct. All rows are counted, regardless of the presence of null values.</li> <li>When tables are being joined, include count_big(*) in the select list to produce the count of the number of rows in the joined results. If the objective is to count the number of rows from one table that match criteria, use count_big(<i>column_name</i>).</li> <li>You can use count_big as an existence check in a subquery. For example: <pre>select * from tab where 0 &lt; (select count_big(*) from tab2 where ...)</pre> <p>However, because count_big counts all matching values, exists or in may return results faster. For example:</p> </li> </ul>

```
select * from tab where exists  
  (select * from tab2 where ...)
```

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute count\_big.

See also

**Commands** compute clause, group by and having clauses, select, where clause

## current\_date

Description Returns the current date.

Syntax `current_date()`

Parameters None.

### Examples

**Example 1** Identifies the current date with `datetime`:

```
1> select datetime(month, current_date())
2> go
```

```
-----
August
```

**Example 2** Identifies the current date with `datepart`:

```
1> select datepart(month, current_date())
2> go
```

```
-----
8
```

(1 row affected)

Usage Finds the current date as it exists on the server.

Standards ANSI SQL – Compliance level: Entry-level compliant.

Permissions Any user can execute `current_date`.

See also **Datatypes** Date and time datatypes

**Commands** `select`, `where` clause

**Functions** `dateadd`, `datetime`, `datepart`, `getdate`

## current\_time

Description Returns the current time.

Syntax current\_time()

Parameters None.

### Examples

**Example 1** Finds the current time:

```
1> select current_date()
2> go
-----
Aug 29 2003

(1 row affected)
```

**Example 2** Use with datetime:

```
1> select datetime(minute, current_time())
2> go
-----
45

(1 row affected)
```

Usage Finds the current time as it exists on the server

Standards ANSI SQL – Compliance level: Entry-level compliant.

Permissions Any user can execute current\_time.

See also **Datatypes** Date and time datatypes

**Commands** select, where clause

**Functions** dateadd, datetime, datepart, getdate



## curunreservedpgs

Description	Returns the number of free pages in the specified disk piece.
Syntax	curunreservedpgs (dbid, lstart, unreservedpgs)
Parameters	<p>dbid is the ID for a database. These are stored in the db_id column of sysdatabases.</p> <p>lstart is a page within the disk piece for which pages are to be returned.</p> <p>unreservedpgs is the default value to return if the dbtable is presently unavailable for the requested database.</p>

**Examples** **Example 1** Returns the database name, device name, and the number of unreserved pages for each device fragment

If a database is open, curunreservedpgs takes the value from memory. If it is not in use, the value is taken from the third parameter you specify in curunreservedpgs. In this example, the value comes from the unreservedpgs column in the sysusages table.

```
select db_name(dbid), d.name,
       curunreservedpgs(dbid, lstart, unreservedpgs)
from sysusages u, sysdevices d
where d.low <= u.size + vstart
      and d.high >= u.size + vstart -1
      and d.status &2 = 2
```

```

                                     name
-----
master                               master                1634
tempdb                               master                423
model                                master                423
pubs2                                 master                 72
sybssystemdb                          master                 399
sybssystemprocs                       master               6577
sybsyntax                             master                 359
```

(7 rows affected)

**Example 2** Displays the number of free pages on the segment for dbid starting on sysusages.lstart:

```
select curunreservedpgs (dbid, sysusages.lstart, 0)
```

Usage	<ul style="list-style-type: none"><li>• <code>curunreservedpgs</code>, a system function, returns the number of free pages in a disk piece. For general information about system functions, see “System functions” on page 68.</li><li>• If a database is open, the value returned by <code>curunreservedpgs</code> is taken from memory. If it is not in use, the value is taken from the third parameter you specify in <code>curunreservedpgs</code>.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>curunreservedpgs</code> .
See also	<b>Functions</b> <code>db_id</code> , <code>lct_admin</code>

## data\_pages

Description	<p>Returns the number of pages used by the specified table, index, or a specific partition. The result does not include pages used for internal structures.</p> <p>This function replaces <code>data_pgs</code> and <code>ptn_data_pgs</code> from versions of Adaptive Server earlier than 15.0.</p>
Syntax	<code>data_pages(dbid, object_id [, indid [, ptnid]])</code>
Parameters	<p><i>dbid</i> is the database ID of the database that contains the data pages.</p> <p><i>object_id</i> is an object ID for a table, view, or other database object. These are stored in the <code>id</code> column of <code>sysobjects</code>.</p> <p><i>indid</i> is the index ID of the target index.</p> <p><i>ptnid</i> is the partition ID of the target partition.</p>
Examples	<p><b>Example 1</b> Returns the number of pages used by the object with a object ID of 31000114 in the specified database (including any indexes):</p> <pre>select data_pages(5, 31000114)</pre> <p><b>Example 2</b> Returns the number of pages used by the object in the data layer, regardless of whether or not a clustered index exists:</p> <pre>select data_pages(5, 31000114, 0)</pre> <p><b>Example 3</b> Returns the number of pages used by the object in the index layer for a clustered index. This does not include the pages used by the data layer:</p> <pre>select data_pages(5, 31000114, 1)</pre> <p><b>Example 4</b> Returns the number of pages used by the object in the data layer of the specific partition, which in this case is 2323242432:</p> <pre>select data_pages(5, 31000114, 0, 2323242432)</pre>
Usage	<p>In the case of an APL (all-pages lock) table, if a clustered index exists on the table, then passing in an <i>indid</i> of:</p> <ul style="list-style-type: none"> <li>• 0 – reports the data pages.</li> <li>• 1 – reports the index pages.</li> </ul> <p>All erroneous conditions return a value of zero, such as when the <i>object_id</i> does not exist in the current database, or the targeted <i>indid</i> or <i>ptnid</i> cannot be found.</p>

Standards                   ANSI SQL – Compliance level: Transact-SQL extension.

Permissions                Any user can execute data\_pages.

See also                    **Functions**   object\_id, row\_count

**System procedure**   sp\_spaceused

## datachange

Description	Measures the amount of change in the data distribution since update statistics last ran. Specifically, it measures the number of inserts, updates, and deletes that have occurred on the given object, partition, or column, and helps you determine if invoking update statistics would benefit the query plan.
Syntax	<code>datachange(object_name, partition_name, column_name)</code>
Parameters	<p><i>object_name</i> is the object name in the current database.</p> <p><i>partition_name</i> is the data partition name. This value can be null.</p> <p><i>column_name</i> is the column name for which the datachange is requested. This value can be null.</p>
Examples	<p><b>Example 1</b> Provides the percentage change in the <code>au_id</code> column in the <code>author_ptn</code> partition:</p> <pre>select datachange("authors", "author_ptn", "au_id")</pre> <p><b>Example 2</b> Provides the percentage change in the <code>authors</code> table on the <code>au_ptn</code> partition. The null value for the <i>column_name</i> parameter indicates that this checks all columns that have histogram statistics and obtains the maximum datachange value from among them.</p> <pre>select datachange("authors", "au_ptn", null)</pre>
Usage	<ul style="list-style-type: none"> <li>• The datachange function requires all three parameters.</li> <li>• datachange is a measure of the inserts, deletes and updates but it does not count them individually. datachange counts an update as a delete and an insert, so each update contributes a count of 2 towards the datachange counter.</li> <li>• The datachange built-in returns the datachange count as a percent of the number of rows, but it bases this percentage on the number of rows remaining, not the original number of rows. For example, if a table has five rows and one row is deleted, datachange reports a value of 25 % since the current row count is 4 and the datachange counter is 1.</li> <li>• datachange is expressed as a percentage of the total number of rows in the table, or partition if you specify a partition. The percentage value can be greater than 100 percent because the number of changes to an object can be much greater than the number of rows in the table, particularly when the number of deletes and updates happening to a table is very high.</li> </ul>

- The value that `datachange` displays is the in-memory value. This can differ from the on-disk value because the on-disk value gets updated by the housekeeper, when you run `sp_flushstats`, or when an object descriptor gets flushed.
- The `datachange` values is not reset when histograms are created for global indexes on partitioned tables.

`datachange` is reset or initialized to zero when:

- New columns are added, and their `datachange` value is initialized.
- New partitions are added, and their `datachange` value is initialized.
- Data-partition-specific histograms are created, deleted or updated. When this occurs, the `datachange` value of the histograms is reset for the corresponding column and partition.
- Data is truncated for a table or partition, and its `datachange` value is reset
- A table is repartitioned either directly or indirectly as a result of some other command, and the `datachange` value is reset for all the table's partitions and columns.
- A table is unpartitioned, and the `datachange` value is reset for all columns for the table.

`datachange` has the following restrictions:

- `datachange` statistics are not maintained on tables in system tempdbs, user-defined tempdbs, system tables, or proxy tables.
- `datachange` updates are non-transactional. If you roll back a transaction, the `datachange` values are not rolled back, and these values can become inaccurate.
- If memory allocation for column-level counters fails, Adaptive Server tracks partition-level `datachange` values instead of column-level values.
- If Adaptive Server does not maintain column-level `datachange` values, it then resets the partition-level `datachange` values whenever the `datachange` values for a column are reset.

Permissions

Any user can execute `datachange`.

## datalength

Description	Returns the actual length, in bytes, of the specified column or string.
Syntax	<code>datalength(expression)</code>
Parameters	<p><i>expression</i></p> <p>is a column name, variable, constant expression, or a combination of any of these that evaluates to a single value. <i>expression</i> can be of any datatype, and is usually a column name. If <i>expression</i> is a character constant, it must be enclosed in quotes.</p>
Examples	<p>Finds the length of the <code>pub_name</code> column in the <code>publishers</code> table:</p> <pre>select Length = datalength(pub_name) from publishers  Length -----       13       16       20</pre>
Usage	<ul style="list-style-type: none"> <li><code>datalength</code>, a system function, returns the length of <i>expression</i> in bytes.</li> <li>For columns defined for the Unicode datatype, <code>datalength</code> returns the actual number of bytes of the data stored in each row. For example, this is what is returned if a <code>unitext</code> column <code>ut</code> contains row value <code>U+0041U+0042U+d800dc00</code>: <pre>select datalength(ut) from unitable -----       8</pre> </li> <li><code>datalength</code> finds the actual length of the data stored in each row. <code>datalength</code> is useful on <code>varchar</code>, <code>univarchar</code>, <code>varbinary</code>, <code>text</code>, and <code>image</code> datatypes, since these datatypes can store variable lengths (and do not store trailing blanks). When a <code>char</code> or <code>unichar</code> value is declared to allow nulls, Adaptive Server stores it internally as <code>varchar</code> or <code>univarchar</code>. For all other datatypes, <code>datalength</code> reports the defined length.</li> <li><code>datalength</code> of any <code>NULL</code> data returns <code>NULL</code>.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>datalength</code> .
See also	<b>Functions</b> <code>char_length</code> , <code>col_length</code>

## dateadd

**Description** Returns the date produced by adding or subtracting a given number of years, quarters, hours, or other date parts to the specified date.

**Syntax** `dateadd(date_part, integer, date expression)`

**Parameters** *date\_part*  
is a date part or abbreviation. For a list of the date parts and abbreviations recognized by Adaptive Server, see “Date parts” on page 64.

*numeric*  
is an integer expression.

*date expression*  
is an expression of type datetime, smalldatetime, date, time, or a character string in a datetime format.

**Examples** **Example 1** Displays the new publication dates when the publication dates of all the books in the titles table slip by 21 days:

```
select newpubdate = dateadd(day, 21, pubdate)
from titles
```

**Example 2** Add one day to a date:

```
declare @a date
select @a = "apr 12, 9999"
select dateadd(dd, 1, @a)
-----
Apr 13 9999
```

**Example 3** Subtracts five minutes to a time:

```
select dateadd(mi, -5, convert(time, "14:20:00"))
-----
2:15PM
```

**Example 4** Add one day to a time and the time remains the same:

```
declare @a time
select @a = "14:20:00"
select dateadd(dd, 1, @a)
-----
2:20PM
```

**Example 5** Although there are limits for each *date\_part*, as with *datetime* values, you can add higher values resulting in the values rolling over to the next significant field:

```
--Add 24 hours to a datetime
```



```
select dateadd(hh, 24, "4/1/1979")
```

```
-----
```

```
Apr  2 1979 12:00AM
```

```
--Add 24 hours to a date
```

```
select dateadd(hh, 24, "4/1/1979")
```

```
-----
```

```
Apr  2 1979
```

## Usage

- `dateadd`, a date function, adds an interval to a specified date. For more information about date functions, see “Date functions” on page 64.
- `dateadd` takes three arguments: the date part, a number, and a date. The result is a `datetime` value equal to the date plus the number of date parts.

If the date argument is a `smalldatetime` value, the result is also a `smalldatetime`. You can use `dateadd` to add seconds or milliseconds to a `smalldatetime`, but such an addition is meaningful only if the result date returned by `dateadd` changes by at least one minute.

- Use the `datetime` datatype only for dates after January 1, 1753. `datetime` values must be enclosed in single or double quotes. Use the `date` datatype for dates from January 1, 0001 to 9999. `date` must be enclosed in single or double quotes. Use `char`, `nchar`, `varchar`, or `nvarchar` for earlier dates. Adaptive Server recognizes a wide variety of date formats. For more information, see “User-defined datatypes” on page 41 and “Datatype conversion functions” on page 55.

Adaptive Server automatically converts between character and `datetime` values when necessary (for example, when you compare a character value to a `datetime` value).

- Using the date part `weekday` or `dw` with `dateadd` is not logical, and produces spurious results. Use `day` or `dd` instead.

**Table 2-7: date\_part recognized abbreviations**

<b>Date part</b>	<b>Abbreviation</b>	<b>Values</b>
Year	yy	1753 – 9999 (datetime) 1900 – 2079 (smalldatetime) 0001 – 9999 (date)
Quarter	qq	1 – 4
Month	mm	1 – 12
Week	wk	1054
Day	dd	1 – 7
dayofyear	dy	1 – 366
Weekday	dw	1 – 7
Hour	hh	0 – 23
Minute	mi	0 – 59
Second	ss	0 – 59
millisecond	ms	0 – 999

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute dateadd.

See also

**Datatypes** Date and time datatypes**Commands** select, where clause**Functions** datediff, datename, datepart, getdate

## datediff

Description	Returns the difference between two dates.
Syntax	<code>datediff(<i>datepart</i>, <i>date expression1</i>, <i>date expression2</i>)</code>
Parameters	<p><i>datepart</i> is a date part or abbreviation. For a list of the date parts and abbreviations recognized by Adaptive Server, see “Date parts” on page 64.</p> <p><i>date expression1</i> is an expression of type datetime, smalldatetime, date, time, or a character string in a datetime format.</p> <p><i>date expression2</i> is an expression of type datetime, smalldatetime, date, time, or a character string in a datetime format.</p>
Examples	<p><b>Example 1</b> Finds the number of days that have elapsed between pubdate and the current date (obtained with the getdate function):</p> <pre>select newdate = datediff(day, pubdate, getdate())       from titles</pre> <p><b>Example 2</b> Find the number of hours between two times:</p> <pre>declare @a time declare @b time select @a = "20:43:22" select @b = "10:43:22" select datediff(hh, @a, @b) ----- -10</pre> <p><b>Example 3</b> Find the number of hours between two dates:</p> <pre>declare @a date declare @b date select @a = "apr 1, 1999" select @b = "apr 2, 1999" select datediff(hh, @a, @b) ----- 24</pre> <p><b>Example 4</b> Find the number of days between two times:</p> <pre>declare @a time declare @b time select @a = "20:43:22" select @b = "10:43:22" select datediff(dd, @a, @b)</pre>

-----  
0

**Example 5** Overflow size of milliseconds return value:

```
select datediff(ms, convert(date, "4/1/1753"), convert(date, "4/1/9999"))
Msg 535, Level 16, State 0:
Line 2:
Difference of two datetime fields caused overflow at runtime.
Command has been aborted
```

**Usage**

- `datediff`, a date function, calculates the number of date parts between two specified dates. For more information about date functions, see “Date functions” on page 64.
- `datediff` takes three arguments. The first is a date part. The second and third are dates. The result is a signed integer value equal to  $date2 - date1$ , in date parts.
- `datediff` produces results of datatype `int`, and causes errors if the result is greater than 2,147,483,647. For milliseconds, this is approximately 24 days, 20:31.846 hours. For seconds, this is 68 years, 19 days, 3:14:07 hours.
- `datediff` results are always truncated, not rounded, when the result is not an even multiple of the date part. For example, using hour as the date part, the difference between “4:00AM” and “5:50AM” is 1.

When you use `day` as the date part, `datediff` counts the number of midnights between the two times specified. For example, the difference between January 1, 1992, 23:00 and January 2, 1992, 01:00 is 1; the difference between January 1, 1992 00:00 and January 1, 1992, 23:59 is 0.

- The month datepart counts the number of first-of-the-months between two dates. For example, the difference between January 25 and February 2 is 1; the difference between January 1 and January 31 is 0.
- When you use the date part `week` with `datediff`, you see the number of Sundays between the two dates, including the second date but not the first. For example, the number of weeks between Sunday, January 4 and Sunday, January 11 is 1.
- If you use `smalldatetime` values, they are converted to `datetime` values internally for the calculation. Seconds and milliseconds in `smalldatetime` values are automatically set to 0 for the purpose of the difference calculation.
- If the second or third argument is a date, and the datepart is hour, minute, second, or millisecond, the dates are treated as midnight.

- If the second or third argument is a time, and the datepart is year, month, or day, then 0 is returned.
- `datediff` results are truncated, not rounded, when the result is not an even multiple of the date part.
- For the smaller time units, there are overflow values, and the function returns an overflow error if you exceed these limits:
  - Milliseconds: approx 24 days
  - Seconds: approx 68 years
  - Minutes: approx 4083 years
  - Others: No overflow limit

Standards ANSI SQL – Compliance level: Transact-SQL extension.

Permissions Any user can execute `datediff`.

See also **Datatypes** Date and time datatypes

**Commands** `select`, `where` clause

**Functions** `dateadd`, `datetime`, `datepart`, `getdate`

## datetime

**Description** Returns the specified datepart (the first argument) of the specified date or time (the second argument) as a character string. Takes a date, time, datetime, or smalldatetime value as its second argument.

**Syntax** datetime (*datepart*, *date expression*)

**Parameters** *datepart*  
 is a date part or abbreviation. For a list of the date parts and abbreviations recognized by Adaptive Server, see “Date parts” on page 64.

*date expression*  
 is an expression of type datetime, smalldatetime, date, time, or a character string in a datetime format.

**Examples** **Example 1** Assumes a current date of November 20, 2000:

```
select datetime(month, getdate())
November
```

**Example 2** Finds the month name of a date:

```
declare @a date
select @a = "apr 12, 0001"
select datetime(mm, @a)
-----
April
```

**Example 3** Finds the seconds of a time:

```
declare @a time
select @a = "20:43:22"
select datetime(ss, @a)
-----
22
```

- Usage**
- datetime, a date function, returns the name of the specified part (such as the month “June”) of a datetime or smalldatetime value, as a character string. If the result is numeric, such as “23” for the day, it is still returned as a character string.
  - For more information about date functions, see “Date functions” on page 64.
  - The date part weekday or dw returns the day of the week (Sunday, Monday, and so on) when used with datetime.
  - Since smalldatetime is accurate only to the minute, when a smalldatetime value is used with datetime, seconds and milliseconds are always 0.

Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute datename.
See also	<b>Datatypes</b> Date and time datatypes <b>Commands</b> select, where clause <b>Functions</b> dateadd, datename, datepart, getdate

## datepart

**Description** Returns the specified datepart in the first argument of the specified date (the second argument) as an integer. Takes a date, time, datetime, or smalldatetime value as its second argument. If the datepart is hour, minute, second, or millisecond, the result is 0.

**Syntax** `datepart(date_part, date expression)`

**Parameters** *date\_part* is a date part. Table 2-8 lists the date parts, the abbreviations recognized by datepart, and the acceptable values.

**Table 2-8: Date parts and their values**

Date part	Abbreviation	Values
year	yy	1753 – 9999 (2079 for smalldatetime). 0001 to 9999 for date
quarter	qq	1 – 4
month	mm	1 – 12
week	wk	1 – 54
day	dd	1 – 31
dayofyear	dy	1 – 366
weekday	dw	1 – 7 (Sun. – Sat.)
hour	hh	0 – 23
minute	mi	0 – 59
second	ss	0 – 59
millisecond	ms	0 – 999
calweekofyear	cwk	1 – 53
calyearofweek	cyr	1753 – 9999 (2079 for smalldatetime). 0001 to 9999 for date
caldayofweek	cdw	1 – 7



When you enter a year as two digits (yy):

- Numbers less than 50 are interpreted as 20yy. For example, 01 is 2001, 32 is 2032, and 49 is 2049.
- Numbers equal to or greater than 50 are interpreted as 19yy. For example, 50 is 1950, 74 is 1974, and 99 is 1999.

Milliseconds can be preceded by either a colon or a period. If preceded by a colon, the number means thousandths of a second. If preceded by a period, a single digit means tenths of a second, two digits mean hundredths of a second, and three digits mean thousandths of a second. For example, “12:30:20:1” means twenty and one-thousandth of a second past 12:30; “12:30:20.1” means twenty and one-tenth of a second past 12:30.

*date expression*

is an expression of type datetime, smalldatetime, date, time, or a character string in a datetime format.

Examples

**Example 1** Assumes a current date of November 25, 1995:

```
select datepart(month, getdate())
-----
11
```

**Example 2** Returns the year of publication from traditional cookbooks:

```
select datepart(year, pubdate) from titles where type =
"trad_cook"
-----
1990
1985
1987
```

**Example 3**

```
select datepart(cwk, '1993/01/01')
-----
53
```

**Example 4**

```
select datepart(cyr, '1993/01/01')
-----
1992
```

**Example 5**

```
select datepart(cdw, '1993/01/01')
-----
5
```

**Example 6** Find the hours in a time:

```
declare @a time
select @a = "20:43:22"
select datepart(hh, @a)
-----
20
```

**Example 7** If a hour, minute, or second portion is requested from a date using datename or datepart) the result is the default time, zero. If a month, day, or year is requested from a time using datename or datepart, the result is the default date, Jan 1 1900:

```
--Find the hours in a date
declare @a date
select @a = "apr 12, 0001"
select datepart(hh, @a)
-----
0

--Find the month of a time
declare @a time
select @a = "20:43:22"
select datename(mm, @a)
-----
January
```

When you give a null value to a datetime function as a parameter, NULL is returned.

## Usage

- datepart, a date function, returns an integer value for the specified part of a datetime value. For more information about date functions, see “Date functions” on page 64.
- datepart returns a number that follows ISO standard 8601, which defines the first day of the week and the first week of the year. Depending on whether the datepart function includes a value for calweekofyear, calyearofweek, or caldayorweek, the date returned may be different for the same unit of time. For example, if Adaptive Server is configured to use U.S. English as the default language, the following returns 1988:

```
datepart(cyr, "1/1/1989")
```

However, the following returns 1989:

```
datepart (yy, "1/1/1989)
```

This disparity occurs because the ISO standard defines the first week of the year as the first week that includes a Thursday *and* begins with Monday.

For servers using U.S. English as their default language, the first day of the week is Sunday, and the first week of the year is the week that contains January 4th.

- The date part `weekday` or `dw` returns the corresponding number when used with `datepart`. The numbers that correspond to the names of weekdays depend on the `datefirst` setting. Some language defaults (including `us_english`) produce `Sunday=1`, `Monday=2`, and so on; others produce `Monday=1`, `Tuesday=2`, and so on. You can change the default behavior on a per-session basis with `set datefirst`. See the `datefirst` option of the `set` command for more information.
- `calweekofyear`, which can be abbreviated as `cwk`, returns the ordinal position of the week within the year. `calyearofweek`, which can be abbreviated as `cyr`, returns the year in which the week begins. `caldayofweek`, which can be abbreviated as `cdw`, returns the ordinal position of the day within the week. You cannot use `calweekofyear`, `calyearofweek`, and `caldayofweek` as date parts for `dateadd`, `datediff`, and `datetime`.
- Since `datetime` and `time` are only accurate to 1/300th of a second, when these datatypes are used with `datepart`, milliseconds are rounded to the nearest 1/300th second.
- Since `smalldatetime` is accurate only to the minute, when a `smalldatetime` value is used with `datepart`, seconds and milliseconds are always 0.
- The values of the weekday date part are affected by the language setting.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute `datepart`.

See also

**Datatypes** Date and time datatypes

**Commands** `select`, `where` clause

**Functions** `dateadd`, `datediff`, `datetime`, `getdate`

## day

Description	Returns an integer that represents the day in the datepart of a specified date.
Syntax	<code>day(date_expression)</code>
Parameters	<i>date_expression</i> is an expression of type datetime, smalldatetime, date, or a character string in a datetime format.
Examples	Returns the integer 02: <pre>day ("11/02/03") ----- 02</pre>
Usage	<code>day(date_expression)</code> is equivalent to <code>datepart(dd,date_expression)</code> .
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute day.
See also	<b>Datatypes</b> datetime, smalldatetime, date, time <b>Functions</b> datepart, month, year

## db\_id

Description	Returns the ID number of the specified database.
Syntax	<code>db_id(database_name)</code>
Parameters	<i>database_name</i> is the name of a database. <i>database_name</i> must be a character expression. If it is a constant expression, it must be enclosed in quotes.
Examples	Returns the ID number of sybssystemprocs: <pre>select db_id("sybssystemprocs") ----- 4</pre>
Usage	<ul style="list-style-type: none"><li>• <code>db_id</code>, a system function, returns the database ID number.</li><li>• If you do not specify a <i>database_name</i>, <code>db_id</code> returns the ID number of the current database.</li><li>• For general information about system functions, see “System functions” on page 68.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>db_id</code> .
See also	<b>Functions</b> <code>db_name</code> , <code>object_id</code>

## db\_name

Description	Returns the name of the database where the ID number is specified.
Syntax	db_name([ <i>database_id</i> ])
Parameters	<i>database_id</i> is a numeric expression for the database ID (stored in sysdatabases.dbid).
Examples	<p><b>Example 1</b> Returns the name of the current database:</p> <pre>select db_name()</pre> <p><b>Example 2</b> Returns the name of database ID 4:</p> <pre>select db_name(4)</pre> <pre>----- sybssystemprocs</pre>
Usage	<ul style="list-style-type: none"><li>• db_name, a system function, returns the database name.</li><li>• If no <i>database_id</i> is supplied, db_name returns the name of the current database.</li><li>• For general information about system functions, see “System functions” on page 68.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute db_name.
See also	<b>Functions</b> col_name, db_id, object_name

## degrees

Description	Returns the size, in degrees, of an angle with the specified number of radians.
Syntax	<code>degrees(<i>numeric</i>)</code>
Parameters	<i>numeric</i> is a number, in radians, to convert to degrees.
Examples	<pre>select degrees(45) -----            2578</pre>
Usage	<ul style="list-style-type: none"> <li>degrees, a mathematical function, converts radians to degrees. Results are of the same type as the numeric expression.</li> </ul> <p>For numeric and decimal expressions, the results have an internal precision of 77 and a scale equal to that of the expression.</p> <ul style="list-style-type: none"> <li>For general information about mathematical functions, see “Mathematical functions” on page 65.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute degrees.
See also	<b>Function</b> radians

## derived\_stat

**Description** Returns derived statistics for the specified object and index.

**Syntax** `derived_stat(object_name | object_id, index_name | index_id, [partition_name | partition_id,] "statistic")`

**Parameters**

*object\_name*  
is the name of the object you are interested in. If you do not specify a fully qualified object name, `derived_stat` searches the current database.

*object\_id*  
is an alternative to *object\_name*, and is the object ID of the object you are interested in. *object\_id* must be in the current database

*index\_name*  
is the name of the index, belonging to the specified object that you are interested in.

*index\_id*  
is an alternative to *index\_name*, and is the index ID of the specified object that you are interested in.

*partition\_name*  
is the name of the partition, belonging to the specific partition that you are interested in.

*partition\_id*  
is an alternative to *partition\_name*, and is the partition ID of the specified object that you are interested in.

"statistic"  
the derived statistic to be returned. Available statistics are:

Value	Returns
data page cluster ratio or dpcr	The data page cluster ratio for the object/index pair
index page cluster ratio or ipcr	The index page cluster ratio for the object/index pair
data row cluster ratio or drcr	The data row cluster ratio for the object/index pair
large io efficiency or lgio	The large I/O efficiency for the object/index pair
space utilization or sput	The space utilization for the object/index pair

**Examples** **Example 1** Selects the space utilization for the titleidind index of the titles table:

```
select derived_stat("titles", "titleidind", "space utilization")
```



**Example 2** Selects the data page cluster ratio for index ID 2 of the titles table. Note that you can use either "dpcr" or "data page cluster ratio":

```
select derived_stat("titles", 2, "dpcr")
```

**Example 3** Statistics are reported for the entire object, as neither the partition ID nor name is not specified:

```
1> select derived_stat(object_id("t1"), 2, "drcr")
2> go
```

```
-----
0.576923
```

**Example 4** Reports the statistic for the partition tl\_928003396:

```
1> select derived_stat(object_id("t1"), 0, "t1_928003306", "drcr")
2> go
```

```
-----
1.000000
```

(1 row affected)

#### Usage

- derived\_stat returns a double precision value.
- The values returned by derived\_stat match the values presented by the optdiag utility.
- If the specified object or index does not exist, derived\_stat returns NULL.
- Specifying an invalid statistic type results in an error message.
- Using the optional *partition\_name* or *partition\_id* reports the target partition; otherwise, derived\_stat reports for the entire object.
- If you provide:
  - Four arguments – derived\_stat uses the third argument as the partition, and returns derived statistics on the fourth argument.
  - Three arguments – derived\_stat assumes you did not specify a partition, and returns derived statistic on the third argument.

#### Standards

ANSI SQL – Compliance level: Transact-SQL extension.

#### Permissions

Only the table owner can execute derived\_stat.

#### See also

**Document** *Performance and Tuning Guide* for:

- “Access Methods and Query Costing for Single Tables”
- “Statistics Tables and Displaying Statistics with optdiag”

**Utility** optdiag

## difference

Description	Returns the difference between two soundex values.
Syntax	<code>difference(<i>expr1</i>,<i>expr2</i>)</code>
Parameters	<p><i>expr1</i> is a character-type column name, variable, or constant expression of char, varchar, nchar, nvarchar, or unichar type.</p> <p><i>expr2</i> is another character-type column name, variable, or constant expression of char, varchar, nchar, nvarchar, or unichar type.</p>
Examples	<p><b>Example 1</b></p> <pre>select difference("smithers", "smothers") ----- 4</pre> <p><b>Example 2</b></p> <pre>select difference("smothers", "brothers") ----- 2</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>difference</code>, a string function, returns an integer representing the difference between two soundex values.</li> <li>• The <code>difference</code> function compares two strings and evaluates the similarity between them, returning a value from 0 to 4. The best match is 4. The string values must be composed of a contiguous sequence of valid single- or double-byte roman letters.</li> <li>• If <i>expr1</i> or <i>expr2</i> is NULL, returns NULL.</li> <li>• If you give a varchar expression as one parameter and a unichar expression as the other, the varchar expression is implicitly converted to unichar (with possible truncation).</li> <li>• For general information about string functions, see “String functions” on page 67.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>difference</code> .
See also	<b>Function</b> <code>soundex</code>

## exp

Description	Returns the value that results from raising the constant to the specified power.
Syntax	<code>exp(<i>approx_numeric</i>)</code>
Parameters	<i>approx_numeric</i> is any approximate numeric (float, real, or double precision) column name, variable, or constant expression.
Examples	<pre>select exp(3) -----                 20.085537</pre>
Usage	<ul style="list-style-type: none"><li>• <code>exp</code>, a mathematical function, returns the exponential value of the specified value.</li><li>• For general information about mathematical functions, see “Mathematical functions” on page 65.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>exp</code> .
See also	<b>Functions</b> <code>log</code> , <code>log10</code> , <code>power</code>

# floor

Description	Returns the largest integer that is less than or equal to the specified value.
Syntax	<code>floor(<i>numeric</i>)</code>
Parameters	<i>numeric</i> is any exact numeric (numeric, dec, decimal, tinyint, smallint, int, or bigint), approximate numeric (float, real, or double precision), or money column, variable, constant expression, or a combination of these.

## Examples

### Example 1

```
select floor(123)
-----
          123
```

### Example 2

```
select floor(123.45)
-----
          123
```

### Example 3

```
select floor(1.2345E2)
-----
        123.000000
```

### Example 4

```
select floor(-123.45)
-----
         -124
```

### Example 5

```
select floor(-1.2345E2)
-----
       -124.000000
```

### Example 6

```
select floor($123.45)
-----
          123.00
```

Usage

- `floor`, a mathematical function, returns the largest integer that is less than or equal to the specified value. Results are of the same type as the numeric expression.

For numeric and decimal expressions, the results have a precision equal to that of the expression and a scale of 0.

- For general information about mathematical functions, see “Mathematical functions” on page 65.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute `floor`.

See also

**Functions** `abs`, `ceiling`, `round`, `sign`

## get\_appcontext

**Description** Returns the value of the attribute in a specified context. `get_appcontext` is a built-in function provided by the Application Context Facility (ACF).

**Syntax** `get_appcontext ("context_name", "attribute_name")`

**Parameters** *context\_name*  
is a row specifying an application context name. It is saved as datatype `char(30)`.

*attribute\_name*  
is a row specifying an application context attribute name. It is saved as datatype `char(30)`.

**Examples** **Example 1** Shows VALUE1 returned for ATTR1.

```
select get_appcontext ("CONTEXT1", "ATTR1")
-----
VALUE1
```

ATTR1 does not exist in CONTEXT2:

```
select get_appcontext ("CONTEXT2", "ATTR1")
```

**Example 2** Shows the result when a user without appropriate permissions attempts to get the application context.

```
select get_appcontext ("CONTEXT1", "ATTR2", "VALUE1")
Select permission denied on built-in get_appcontext, database dbid
-----
-1
```

**Usage**

- This function returns 0 for success and -1 for failure.
- If the attribute you require does not exist in the application context, `get_appcontext` returns NULL.
- `get_appcontext` saves attributes as `char` datatypes. If you are creating an access rule that compares the attribute value to other datatypes, the rule should convert the `char` data to the appropriate datatype.
- All arguments for this function are required.

**Standards** ANSI SQL – Compliance level: Transact-SQL extension.

**Permissions** Permissions depend on the user profile and the application profile, and are stored by the ACF.

**See also** For more information on the ACF, see “Row-level access control” in Chapter 11, “Managing User Permissions” of the *System Administration Guide*.

**Functions** get\_appcontext, list\_appcontext, rm\_appcontext, set\_appcontext



## getdate

Description	Returns the current system date and time.
Syntax	getdate()
Parameters	None.
Examples	<p><b>Example 1</b> Assumes a current date of November 25, 1995, 10:32 a.m.:</p> <pre>select getdate() Nov 25 1995 10:32AM</pre> <p><b>Example 2</b> Assumes a current date of November:</p> <pre>select datepart(month, getdate()) 11</pre> <p><b>Example 3</b> Assumes a current date of November:</p> <pre>select datename(month, getdate()) November</pre>
Usage	<ul style="list-style-type: none"><li>• getdate, a date function, returns the current system date and time.</li><li>• For more information about date functions, see “Date functions” on page 64.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute getdate.
See also	<p><b>Datatypes</b> Date and time datatypes</p> <p><b>Functions</b> dateadd, datediff, datename, datepart</p>

## **getutcdate**

Description	Returns a date and time where the value is in Universal Coordinated Time (UTC). <i>getutcdate</i> is calculated each time a row is inserted or selected.
Syntax	<code>insert t1 (c1, c2, c3) select c1, getutcdate(), getdate() from t2)</code>
Usage	Returns a date and time that has a value in Universal Coordinated Time (UTC). <i>getutcdate</i> is calculated each time a row is inserted or selected.
See also	<b>Functions</b> <code>biginttohex</code> , <code>convert</code>

## has\_role

Description	Returns information about whether the user has been granted the specified role.
Syntax	<code>has_role ("role_name"[, 0])</code>
Parameters	<p><i>role_name</i> is the name of a system or user-defined role.</p> <p>0 is an optional parameter that suppresses auditing.</p>
Examples	<p><b>Example 1</b> Creates a procedure to check if the user is a System Administrator:</p> <pre>create procedure sa_check as if (has_role("sa_role", 0) &gt; 0) begin     print "You are a System Administrator."     return(1) end</pre> <p><b>Example 2</b> Checks that the user has been granted the System Security Officer role:</p> <pre>select has_role("sso_role", 0)</pre> <p><b>Example 3</b> Checks that the user has been granted the Operator role:</p> <pre>select has_role("oper_role", 0)</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>has_role</code> functions the same way <code>proc_role</code> does. Beginning with Adaptive Server version 15.0, Sybase supports—and recommends—that you use <code>has_role</code> instead of <code>proc_role</code>. You need not, however, convert all of your existing uses of <code>proc_role</code> to <code>has_role</code>.</li> <li>• <code>has_role</code>, a system function, checks whether an invoking user has been granted, and has activated, the specified role.</li> <li>• <code>has_role</code> returns 0 if the user has: <ul style="list-style-type: none"> <li>• Not been granted the specified role</li> <li>• Not been granted a role which contains the specified role</li> <li>• Been granted, but has not activated, the specified role</li> </ul> </li> <li>• <code>has_role</code> returns 1 if the invoking user has been granted, and has activated, the specified role.</li> <li>• <code>has_role</code> returns 2 if the invoking user has a currently active role, which contains the specified role.</li> </ul>

- For general information about system functions, see “System functions” on page 68.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute `has_role`.

See also

**Commands** `alter role`, `create role`, `drop role`, `grant`, `set`, `revoke`

**Functions** `mut_excl_roles`, `role_contain`, `role_id`, `role_name`, `show_role`

## hextobigint

Description	Returns the bigint value equivalent of a hexadecimal string
Syntax	<code>hextobigint (hexadecimal_string)</code>
Parameters	<p><i>hexadecimal_string</i></p> <p>is the hexadecimal value to be converted to an big integer; must be a character-type column, variable name, or a valid hexadecimal string, with or without a “0x” prefix, enclosed in quotes.</p>
Examples	<p>The following example converts the hexadecimal string 0x7fffffffffffffff to a big integer.</p> <pre> 1&gt; select hextobigint ("0x7fffffffffffffff") 2&gt; go ----- 9223372036854775807 </pre>
Usage	<ul style="list-style-type: none"> <li>• <code>hextobigint</code>, a datatype conversion function, returns the platform-independent integer equivalent of a hexadecimal string.</li> <li>• Use the <code>hextobigint</code> function for platform-independent conversions of hexadecimal data to integers. <code>hextobigint</code> accepts a valid hexadecimal string, with or without a “0x” prefix, enclosed in quotes, or the name of a character-type column or variable.</li> </ul> <p><code>hextobigint</code> returns the bigint equivalent of the hexadecimal string. The function always returns the same bigint equivalent for a given hexadecimal string, regardless of the platform on which it is executed.</p>
See also	<b>Functions</b> <code>biginttohex</code> , <code>convert</code> , <code>inttohex</code> , <code>hextoint</code>

## hextoint

Description	Returns the platform-independent integer equivalent of a hexadecimal string.
Syntax	hextoint ( <i>hexadecimal_string</i> )
Parameters	<i>hexadecimal_string</i> is the hexadecimal value to be converted to an integer; must be a character-type column, variable name, or a valid hexadecimal string, with or without a “0x” prefix, enclosed in quotes.
Examples	Returns the integer equivalent of the hexadecimal string “0x00000100”. The result is always 256, regardless of the platform on which it is executed: <pre>select hextoint ("0x00000100")</pre>
Usage	<ul style="list-style-type: none"><li>• hextoint, a datatype conversion function, returns the platform-independent integer equivalent of a hexadecimal string.</li><li>• Use the hextoint function for platform-independent conversions of hexadecimal data to integers. hextoint accepts a valid hexadecimal string, with or without a “0x” prefix, enclosed in quotes, or the name of a character-type column or variable.  hextoint returns the integer equivalent of the hexadecimal string. The function always returns the same integer equivalent for a given hexadecimal string, regardless of the platform on which it is executed.</li><li>• For more information about datatype conversion, see “Datatype conversion functions” on page 55.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute hextoint.
See also	<b>Functions</b> biginttohex, convert, intohex

## host\_id

Description	Returns the client computer's operating system process ID for the current Adaptive Server client.
Syntax	host_id()
Parameters	None.
Examples	<p>In this example, the name of the client computer is "ephemeris" and the process ID on the computer "ephemeris" for the Adaptive Server client process is 2309:</p> <pre>select host_name(), host_id() ----- ephemeris                2309</pre> <p>The following is the process information, gathered using the UNIX ps command, from the computer "ephemeris" showing that the client in this example is "isql" and its process ID is 2309:</p> <pre>2309 pts/2    S   0:00 /work/as125/OCS-12_5/bin/isql</pre>
Usage	<ul style="list-style-type: none"> <li>• host_id, a system function, returns the host process ID of the client process (not the server process).</li> <li>• For general information about system functions, see "String functions" on page 67.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute host_id.
See also	<b>Function</b> host_name

## host\_name

Description Returns the current host computer name of the client process.

Syntax host\_name()

Parameters None.

Examples 

```
select host_name()  
-----  
violet
```

Usage

- host\_name, a system function, returns the current host computer name of the client process (not the server process).
- For general information about system functions, see “System functions” on page 68.

Standards ANSI SQL – Compliance level: Transact-SQL extension.

Permissions Any user can execute host\_name.

See also **Function** host\_id



## identity\_burn\_max

Description	Tracks the identity burn max value for a given table. This function returns only the value; does not perform an update.
Syntax	<code>identity_burn_max(<i>table_name</i>)</code>
Parameters	<i>table_name</i> is the name of the table selected.
Examples	<pre>select identity_burn_max("t1") t1 ----- 51</pre>
Usage	<code>identity_burn_max</code> tracks the identity burn max value for a given table.
Permissions	Only the table owner, System Administrator, or database administrator can issue this command.

## index\_col

Description	Returns the name of the indexed column in the specified table or view, and can be up to 255 bytes in length
Syntax	<code>index_col (object_name, index_id, key_#[, user_id])</code>
Parameters	<p><i>object_name</i> is the name of a table or view. The name can be fully qualified (that is, it can include the database and owner name). It must be enclosed in quotes.</p> <p><i>index_id</i> is the number of <i>object_name</i>'s index. This number is the same as the value of <code>sysindexes.indid</code>.</p> <p><i>key_#</i> is a key in the index. This value is between 1 and <code>sysindexes.keycnt</code> for a clustered index and between 1 and <code>sysindexes.keycnt+1</code> for a nonclustered index.</p> <p><i>user_id</i> is the owner of <i>object_name</i>. If you do not specify <i>user_id</i>, it defaults to the caller's user ID.</p>
Examples	<p>Finds the names of the keys in the clustered index on table t4:</p> <pre>declare @keycnt integer select @keycnt = keycnt from sysindexes       where id = object_id("t4")       and indid = 1 while @keycnt &gt; 0 begin     select index_col("t4", 1, @keycnt)     select @keycnt = @keycnt - 1 end</pre>
Usage	<ul style="list-style-type: none"><li>• <code>index_col</code>, a system function, returns the name of the indexed column.</li><li>• <code>index_col</code> returns NULL if <i>object_name</i> is not a table or view name.</li><li>• For general information about system functions, see “String functions” on page 67.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>index_col</code> .
See also	<b>Function</b> <code>object_id</code> <b>System procedure</b> <code>sp_helpindex</code>

## index\_colorder

Description	Returns the column order.
Syntax	<code>index_colorder (object_name, index_id, key_# [, user_id])</code>
Parameters	<p><i>object_name</i> is the name of a table or view. The name can be fully qualified (that is, it can include the database and owner name). It must be enclosed in quotes.</p> <p><i>index_id</i> is the number of <i>object_name</i>'s index. This number is the same as the value of <code>sysindexes.indid</code>.</p> <p><i>key_#</i> is a key in the index. Valid values are 1 and the number of keys in the index. The number of keys is stored in <code>sysindexes.keycnt</code>.</p> <p><i>user_id</i> is the owner of <i>object_name</i>. If you do not specify <i>user_id</i>, it defaults to the caller's user ID.</p>
Examples	<p>Returns "DESC" because the salesind index on the sales table is in descending order:</p> <pre>select name, index_colorder("sales", indid, 2) from sysindexes where id = object_id ("sales") and indid &gt; 0  name ----- salesind          DESC</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>index_colorder</code>, a system function, returns "ASC" for columns in ascending order or "DESC" for columns in descending order.</li> <li>• <code>index_colorder</code> returns NULL if <i>object_name</i> is not a table name or if <i>key_#</i> is not a valid key number.</li> <li>• For general information about system functions, see "String functions" on page 67.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>index_colorder</code> .

## inttohex

Description	Returns the platform-independent hexadecimal equivalent of the specified integer.
Syntax	<code>inttohex (integer_expression)</code>
Parameters	<i>integer_expression</i> is the integer value to be converted to a hexadecimal string.
Examples	<pre>select inttohex (10) ----- 0000000A</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>inttohex</code>, a datatype conversion function, returns the platform-independent hexadecimal equivalent of an integer, without a “0x” prefix.</li> <li>• Use the <code>inttohex</code> function for platform-independent conversions of integers to hexadecimal strings. <code>inttohex</code> accepts any expression that evaluates to an integer. It always returns the same hexadecimal equivalent for a given expression, regardless of the platform on which it is executed.</li> <li>• For more information about datatype conversion, see “Datatype conversion functions” on page 55.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>inttohex</code> .
See also	<b>Functions</b> <code>convert</code> , <code>hextobigint</code> , <code>hextoint</code>

## is\_quiesced

Description	Indicates whether a database is in quiesce database mode. <code>is_quiesced</code> returns 1 if the database is quiesced and 0 if it is not.
Syntax	<code>is_quiesced(dbid)</code>
Parameters	<i>dbid</i> is the database ID of the database.
Examples	<b>Example 1</b> Uses the test database, which has a database ID of 4, and which is not quiesced:

```
1> select is_quiesced(4)
2> go
```

```
-----
                0
```

(1 row affected)

**Example 2** Uses the test database after running quiesce database to suspend activity:

```
1> quiesce database tst hold test
2> go
1> select is_quiesced(4)
2> go
```

```
-----
                1
```

(1 row affected)

**Example 3** Uses the test database after resuming activity using quiesce database:

```
1> quiesce database tst release
2> go
1> select is_quiesced(4)
2> go
```

```
-----
                0
```

(1 row affected)

**Example 4** Executes a select statement with `is_quiesced` using an invalid database ID:

```
1>select is_quiesced(-5)
```

```
2> go
-----
      NULL

(1 row affected)
```

Usage

- `is_quiesced` has no default values. You see an error if you execute `is_quiesced` without specifying a database.
- `is_quiesced` returns NULL if you specify a database ID that does not exist.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute `is_quiesced`.

See also

**Command** `quiesce database`

## is\_sec\_service\_on

Description	Returns 1 if the security service is active and 0 if it is not.
Syntax	<code>is_sec_service_on(<i>security_service_nm</i>)</code>
Parameters	<i>security_service_nm</i> is the name of the security service.
Examples	<pre>select is_sec_service_on("unifiedlogin")</pre>
Usage	<ul style="list-style-type: none"> <li>Use <code>is_sec_service_on</code> to determine whether a given security service is active during the session.</li> <li>To find valid names of security services, execute: <pre>select * from syssecmechs</pre> <p>The result might look something like:</p> <pre>sec_mech_name  available_service ----- dce            unifiedlogin dce            mutualauth dce            delegation dce            integrity dce            confidentiality dce            detectreplay dce            detectseq</pre> <p>The <code>available_service</code> column displays the security services that are supported by Adaptive Server.</p> </li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>is_sec_service_on</code> .
See also	<b>Function</b> <code>show_sec_services</code>

## isnull

Description	Substitutes the value specified in <i>expression2</i> when <i>expression1</i> evaluates to NULL.
Syntax	<code>isnull(<i>expression1</i>, <i>expression2</i>)</code>
Parameters	<i>expression</i> is a column name, variable, constant expression, or a combination of any of these that evaluates to a single value. It can be of any datatype, including <code>unichar</code> . <i>expression</i> is usually a column name. If <i>expression</i> is a character constant, it must be enclosed in quotes.
Examples	Returns all rows from the titles table, replacing null values in price with 0: <pre>select isnull(price,0) from titles</pre>
Usage	<ul style="list-style-type: none"><li>• <code>isnull</code>, a system function, substitutes the value specified in <i>expression2</i> when <i>expression1</i> evaluates to NULL. For general information about system functions, see “String functions” on page 67.</li><li>• The datatypes of the expressions must convert implicitly, or you must use the <code>convert</code> function.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>isnull</code> .
See also	<b>Function</b> <code>convert</code>



## lct\_admin

Description	Manages the last-chance threshold, returns the current value of the last-chance threshold (LCT), and aborts transactions in a transaction log that has reached its LCT.
Syntax	<pre>lct_admin({"lastchance"   "logfull"   "reserved_for_rollbacks"},          database_id            "reserve", {log_pages   0 }            "abort", process-id [, database-id])</pre>
Parameters	<p><b>lastchance</b> creates a LCT in the specified database.</p> <p><b>logfull</b> returns 1 if the LCT has been crossed in the specified database and 0 if it has not.</p> <p><b>reserved_for_rollbacks</b> determines the number of pages a database currently reserved for rollbacks.</p> <p><b>database_id</b> specifies the database.</p> <p><b>reserve</b> obtains either the current value of the LCT or the number of log pages required for dumping a transaction log of a specified size.</p> <p><b>log_pages</b> is the number of pages for which to determine a LCT.</p> <p><b>0</b> returns the current value of the LCT. The size of the LCT in a database with separate log and data segments does not vary dynamically. It has a fixed value, based on the size of the transaction log. The LCT varies dynamically in a database with mixed log and data segments.</p> <p><b>abort</b> aborts transactions in a database where the transaction log has reached its last-chance threshold. Only transactions in log-suspend mode can be aborted.</p> <p><b>logsegment_freepages</b> describes the free space available for the log segment. This is the total value of free space, not per-disk.</p>

*process-id*

The ID (*spid*) of a process in log-suspend mode. A process is placed in log-suspend mode when it has open transactions in a transaction log that has reached its last-chance threshold (LCT).

*database-id*

the ID of a database with a transaction log that has reached its LCT. If *process-id* is 0, all open transactions in the specified database are terminated.

## Examples

**Example 1** Creates the log segment last-chance threshold for the database with dbid 1. It returns the number of pages at which the new threshold resides. If there was a previous last-chance threshold, it is replaced:

```
select lct_admin("lastchance", 1)
```

**Example 2** Returns 1 if the last-chance threshold for the database with dbid of 6 has been crossed, and 0 if it has not:

```
select lct_admin("logfull", 6)
```

**Example 3** Calculates and returns the number of log pages that would be required to successfully dump the transaction log in a log containing 64 pages:

```
select lct_admin("reserve", 64)
```

```
-----  
16
```

**Example 4** Returns the current last-chance threshold of the transaction log in the database from which the command was issued:

```
select lct_admin("reserve", 0)
```

**Example 5** Aborts transactions belonging to process 83. The process must be in log-suspend mode. Only transactions in a transaction log that has reached its LCT are terminated:

```
select lct_admin("abort", 83)
```

**Example 6** Aborts all open transactions in the database with dbid of 5. This form awakens any processes that may be suspended at the log segment last-chance threshold:

```
select lct_admin("abort", 0, 5)
```

**Example 7** Determines the number of pages reserved for rollbacks in the pubs2 database, which has a dbid of 5:

```
select lct_admin("reserved_for_rollbacks", 5, 0)
```

**Example 8** Describes the free space available for a database with a dbid of 4:

---

Usage	<pre>select lct_admin("logsegment_freepages", 4)</pre> <ul style="list-style-type: none"> <li>• <code>lct_admin</code>, a system function, manages the log segment's last-chance threshold. For general information about system functions, see "System functions" on page 68.</li> <li>• If <code>lct_admin("lastchance", <i>dbid</i>)</code> returns zero, the log is not on a separate segment in this database, so no last-chance threshold exists.</li> <li>• Whenever you create a database with a separate log segment, the server creates a default last chance threshold that defaults to calling <code>sp_thresholdaction</code>. This happens even if a procedure called <code>sp_thresholdaction</code> does not exist on the server at all.</li> </ul> <p>If your log crosses the last-chance threshold, Adaptive Server suspends activity, tries to call <code>sp_thresholdaction</code>, finds it does not exist, generates an error, then leaves processes suspended until the log can be truncated.</p> <ul style="list-style-type: none"> <li>• To terminate the oldest open transaction in a transaction log that has reached its LCT, enter the ID of the process that initiated the transaction.</li> <li>• To terminate all open transactions in a transaction log that has reached its LCT, enter 0 as the <i>process-id</i>, and specify a database ID in the <i>database-id</i> parameter.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Only a System Administrator can execute <code>lct_admin abort</code> . Any user can execute the other <code>lct_admin</code> options.
See also	<p><b>Document</b> <i>System Administration Guide</i>.</p> <p><b>Command</b> <code>dump transaction</code></p> <p><b>Function</b> <code>curunreservedpgs</code></p> <p><b>System procedure</b> <code>sp_thresholdaction</code></p>

## left

**Description** Returns a specified number of characters on the left end of a character string.

**Syntax** `left(character_expression, integer_expression)`

**Parameters** *character\_expression*  
is the character string from which the characters on the left are selected.

*integer\_expression*  
is the positive integer that specifies the number of characters returned. An error is returned if *integer\_expression* is negative.

**Examples** **Example 1** Returns the five leftmost characters of each book title.

```
use pubs
select left(title, 5)
from titles
order by title_id
```

```
-----
The B
Cooki
You C
.....
Sushi
```

```
(18 row(s) affected)
```

**Example 2** Returns the two leftmost characters of the character string "abcdef".

```
select left("abcdef", 2)
-----
ab
(1 row(s) affected)
```

**Usage**

- *character\_expression* can be of any datatype (except text or image) that can be implicitly converted to varchar or nvarchar. *character\_expression* can be a constant, variable, or a column name. You can explicitly convert *character\_expression* using convert.
- left is equivalent to substring(*character\_expression*, 1, *integer\_expression*). For more information on this function, see substring on page 238.

**Standards** ANSI SQL – Compliance level: Transact-SQL extension.

**Permissions** Any user can execute left.

**See also** **Datatypes** varchar, nvarchar

**Functions** len, str\_replace, substring

## len

Description	Returns the number of characters, not the number of bytes, of a specified string expression, excluding trailing blanks.
Syntax	<code>len(<i>string_expression</i>)</code>
Parameters	<i>string_expression</i> is the string expression to be evaluated.
Examples	Returns the characters <pre>select len(notes) from titles where title_id = "PC9999" ----- 39</pre>
Usage	This function is the equivalent of <code>char_length(<i>string_expression</i>)</code> .
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute len.
See also	<b>Datatypes</b> char, nchar, varchar, nvarchar <b>Functions</b> char_length, left, str_replace

## license\_enabled

Description	Returns 1 if a feature's license is enabled, 0 if the license is not enabled, or NULL if you specify an invalid license name.
Syntax	<code>license_enabled("ase_server"   "ase_ha"   "ase_dtm"   "ase_java"   "ase_asm")</code>
Parameters	<p><code>ase_server</code> specifies the license for Adaptive Server.</p> <p><code>ase_ha</code> specifies the license for the Adaptive Server high availability feature.</p> <p><code>ase_dtm</code> specifies the license for Adaptive Server distributed transaction management features.</p> <p><code>ase_java</code> specifies the license for the Java in Adaptive Server feature.</p> <p><code>ase_asm</code> specifies the license for Adaptive Server advanced security mechanism.</p>
Examples	<p>Indicates that the license for the Adaptive Server distributed transaction management feature is enabled:</p> <pre> select license_enabled("ase_dtm") ----- 1 </pre>
Usage	<ul style="list-style-type: none"> <li>For information about installing license keys for Adaptive Server features, see your installation guide.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>license_enabled</code> .
See also	<p><b>Documents</b> Installation guide for your platform</p> <p><b>System procedure</b> <code>sp_configure</code></p>

## list\_appcontext

**Description** Lists all the attributes of all the contexts in the current session. list\_appcontext is a built-in function provided by the Application Context Facility (ACF).

**Syntax** list\_appcontext (["context\_name"])

**Parameters** *context\_name*  
is an optional argument that names all the application context attributes in the session.

**Examples** **Example 1** Shows the results when a user with appropriate permissions attempts to list the application contexts:

```
select list_appcontext ([context_name])  
  
Context Name: (CONTEXT1)  
Attribute Name: (ATTR1) Value: (VALUE2)  
Context Name: (CONTEXT2)  
Attribute Name: (ATTR1) Value: (VALUE1)
```

**Example 2** Shows the results when a user without appropriate permissions attempts to list the application contexts:

```
select list_appcontext()  
  
Select permission denied on built-in list_appcontext,  
database DBID  
-----  
-1
```

**Usage**

- This function returns 0 for success.
- Since built-in functions do not return multiple result sets, the client application receives list\_appcontext returns as messages.

**Standards** ANSI SQL – Compliance level: Transact-SQL extension

**Permissions** Permissions depend on the user profile and the application profile, and are stored by the ACF.

**See also** For more information on the ACF, see “Row-level access control” in Chapter 11, “Managing User Permissions” of the *System Administration Guide*.

**Functions** get\_appcontext, list\_appcontext, rm\_appcontext, set\_appcontext



## lockscheme

Description	Returns the locking scheme of the specified object as a string.
Syntax	<pre>lockscheme(<i>object_name</i>) lockscheme(<i>object_id</i> [, <i>db_id</i>])</pre>
Parameters	<p><i>object_name</i> is the name of the object that the locking scheme returns. <i>object_name</i> can also be a fully qualified name.</p> <p><i>db_id</i> the ID of the database specified by <i>object_id</i>.</p> <p><i>object_id</i> the ID of the object that the locking scheme returns.</p>
Examples	<p><b>Example 1</b> Selects the locking scheme for the titles table in the current database:</p> <pre>select lockscheme("titles")</pre> <p><b>Example 2</b> Selects the locking scheme for <i>object_id</i> 224000798 (in this case, the titles table) from database ID 4 (the pubs2 database):</p> <pre>select lockscheme(224000798, 4)</pre> <p><b>Example 3</b> Returns the locking scheme for the titles table (<i>object_name</i> in this example is fully qualified):</p> <pre>select lockscheme(tempdb.ownerjoe.titles)</pre>
Usage	<ul style="list-style-type: none"> <li>lockscheme returns varchar(11) and allows NULLs.</li> <li>lockscheme defaults to the current database if you: <ul style="list-style-type: none"> <li>Do not provide a fully qualified <i>object_name</i>.</li> <li>Do not provide a <i>db_id</i>.</li> <li>Provide a null for <i>db_id</i>.</li> </ul> </li> <li>If the specified object is not a table, lockscheme returns the string “not a table.”</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute lockscheme.

## log

Description	Returns the natural logarithm of the specified number.
Syntax	<code>log(<i>approx_numeric</i>)</code>
Parameters	<i>approx_numeric</i> is any approximate numeric (float, real, or double precision) column name, variable, or constant expression.
Examples	<pre>select log(20) -----                 2.995732</pre>
Usage	<ul style="list-style-type: none"><li>• log, a mathematical function, returns the natural logarithm of the specified value.</li><li>• For general information about mathematical functions, see “Mathematical functions” on page 65.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute log.
See also	<b>Functions</b> log10, power

## log10

Description	Returns the base 10 logarithm of the specified number.
Syntax	<code>log10(<i>approx_numeric</i>)</code>
Parameters	<i>approx_numeric</i> is any approximate numeric (float, real, or double precision) column name, variable, or constant expression.
Examples	<pre>select log10(20) -----            1.301030</pre>
Usage	<ul style="list-style-type: none"><li>• <code>log10</code>, a mathematical function, returns the base 10 logarithm of the specified value.</li><li>• For general information about mathematical functions, see “Mathematical functions” on page 65.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>log10</code> .
See also	<b>Functions</b> <code>log</code> , <code>power</code>

## lower

Description	Returns the lowercase equivalent of the specified expression.
Syntax	<code>lower(char_expr   uchar_expr)</code>
Parameters	<p><i>char_expr</i> is a character-type column name, variable, or constant expression of char, varchar, nchar, or nvarchar type.</p> <p><i>uchar_expr</i> is a character-type column name, variable, or constant expression of unichar or univarchar type.</p>
Examples	<pre>select lower(city) from publishers ----- boston washington berkeley</pre>
Usage	<ul style="list-style-type: none"><li>• lower, a string function, converts uppercase to lowercase, returning a character value.</li><li>• lower is the inverse of upper.</li><li>• If <i>char_expr</i> or <i>uchar_expr</i> is NULL, returns NULL.</li><li>• For general information about string functions, see “String functions” on page 67.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute lower.
See also	<b>Function</b> upper

## ltrim

Description	Returns the specified expression, trimmed of leading blanks.
Syntax	<code>ltrim(char_expr   uchar_expr)</code>
Parameters	<p><i>char_expr</i> is a character-type column name, variable, or constant expression of char, varchar, nchar, or nvarchar type.</p> <p><i>uchar_expr</i> is a character-type column name, variable, or constant expression of unichar or univarchar type.</p>
Examples	<pre>select ltrim("  123") ----- 123</pre>
Usage	<ul style="list-style-type: none"> <li>• ltrim, a string function, removes leading blanks from the character expression. Only values equivalent to the space character in the current character set are removed.</li> <li>• If <i>char_expr</i> or <i>uchar_expr</i> is NULL, returns NULL.</li> <li>• For Unicode expressions, returns the lowercase Unicode equivalent of the specified expression. Characters in the expression that have no lowercase equivalent are left unmodified.</li> <li>• For general information about string functions, see “String functions” on page 67.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute ltrim.
See also	<b>Function</b> rtrim

## max

Description	Returns the highest value in an expression.
Syntax	<code>max(<i>expression</i>)</code>
Parameters	<i>expression</i> is a column name, constant, function, any combination of column names, constants, and functions connected by arithmetic or bitwise operators, or a subquery.
Examples	<p><b>Example 1</b> Returns the maximum value in the discount column of the salesdetail table as a new column:</p> <pre>select max(discount) from salesdetail -----                 62.200000</pre> <p><b>Example 2</b> Returns the maximum value in the discount column of the salesdetail table as a new row:</p> <pre>select discount from salesdetail compute max(discount)</pre>
Usage	<ul style="list-style-type: none"><li>• max, an aggregate function, finds the maximum value in a column or expression. For general information about aggregate functions, see “Aggregate functions” on page 49.</li><li>• You can use max with exact and approximate numeric, character, and datetime columns; you cannot use it with bit columns. With character columns, max finds the highest value in the collating sequence. max ignores null values. max implicitly converts char datatypes to varchar, and unichar datatypes to univarchar, stripping all trailing blanks.</li><li>• unichar data is collated according to the default Unicode sort order.</li><li>• Adaptive Server goes directly to the end of the index to find the last row for max when there is an index on the aggregated column, unless:<ul style="list-style-type: none"><li>• The <i>expression</i> not a column.</li><li>• The column is not the first column of an index.</li><li>• There is another aggregate in the query.</li><li>• There is a group by or where clause.</li></ul></li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute max.

See also

**Commands** compute clause, group by and having clauses, select, where clause

**Functions** avg, min

## min

Description	Returns the lowest value in a column.
Syntax	<code>min(expression)</code>
Parameters	<p><i>expression</i></p> <p>is a column name, constant, function, any combination of column names, constants, and functions connected by arithmetic or bitwise operators, or a subquery. With aggregates, an expression is usually a column name. For more information, see “Expressions” on page 275.</p>
Examples	<pre>select min(price) from titles       where type = "psychology" -----                         7.00</pre>
Usage	<ul style="list-style-type: none"><li>• min, an aggregate function, finds the minimum value in a column.</li><li>• For general information about aggregate functions, see “Aggregate functions” on page 49.</li><li>• You can use min with numeric, character, time, and datetime columns; you cannot use it with bit columns. With character columns, min finds the lowest value in the sort sequence. min implicitly converts char datatypes to varchar, and unichar datatypes to univarchar, stripping all trailing blanks. min ignores null values. distinct is not available, since it is not meaningful with min.</li><li>• unichar data is collated according to the default Unicode sort order.</li><li>• Adaptive Server goes directly to the first qualifying row for min when there is an index on the aggregated column, unless:<ul style="list-style-type: none"><li>• The <i>expression</i> is not a column.</li><li>• The column is not the first column of an index.</li><li>• There is another aggregate in the query.</li><li>• There is a group by clause.</li></ul></li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute min.
See also	<p><b>Commands</b> compute clause, group by and having clauses, select, where clause</p> <p><b>Functions</b> avg, max</p>



## month

Description	Returns an integer that represents the month in the datepart of a specified date.
Syntax	<code>month(date_expression)</code>
Parameters	<i>date_expression</i> is an expression of type <code>datetime</code> , <code>smalldatetime</code> , <code>date</code> , or a character string in a <code>datetime</code> format.
Examples	Returns the integer 11: <pre>day ("11/02/03") ----- 11</pre>
Usage	<code>month(date_expression)</code> is equivalent to <code>datepart(mm, date_expression)</code> .
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>month</code> .
See also	<b>Datatypes</b> <code>datetime</code> , <code>smalldatetime</code> , <code>date</code> <b>Functions</b> <code>datepart</code> , <code>day</code> , <code>year</code>

## mut\_excl\_roles

Description	Returns information about the mutual exclusivity between two roles.
Syntax	<code>mut_excl_roles (role1, role2 [membership   activation])</code>
Parameters	<p><i>role1</i> is one user-defined role in a mutually exclusive relationship.</p> <p><i>role2</i> is the other user-defined role in a mutually exclusive relationship.</p> <p><i>level</i> is the level (membership or activation) at which the specified roles are exclusive.</p>
Examples	<p>Shows that the admin and supervisor roles are mutually exclusive:</p> <pre>alter role admin add exclusive membership supervisor select mut_excl_roles("admin", "supervisor", "membership") ----- 1</pre>
Usage	<ul style="list-style-type: none"><li>• <code>mut_excl_roles</code>, a system function, returns information about the mutual exclusivity between two roles. If the System Security Officer defines <code>role1</code> as mutually exclusive with <code>role2</code> or a role directly contained by <code>role2</code>, <code>mut_excl_roles</code> returns 1. If the roles are not mutually exclusive, <code>mut_excl_roles</code> returns 0.</li><li>• For general information about system functions, see “System functions” on page 68.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension
Permissions	Any user can execute <code>mut_excl_roles</code> .
See also	<p><b>Commands</b> alter role, create role, drop role, grant, set, revoke</p> <p><b>Functions</b> <code>proc_role</code>, <code>role_contain</code>, <code>role_id</code>, <code>role_name</code></p> <p><b>System procedures</b> <code>sp_activeroles</code>, <code>sp_displayroles</code>, <code>sp_role</code></p>

## newid

Description	Generates human-readable, globally unique IDs (GUIDs) in two different formats, based on arguments you provide. The length of the human-readable format of the GUID value is either 32 bytes (with no dashes) or 36 bytes (with dashes).
Syntax	<code>newid([optionflag])</code>
Parameters	<p><i>option flag</i></p> <ul style="list-style-type: none"> <li>• 0, or no value – the GUID generated is human-readable (varchar), but does not include dashes. This argument, which is the default, is useful for converting values into varbinary.</li> <li>• -1 – the GUID generated is human-readable (varchar) and includes dashes.</li> <li>• -0x0 – returns the GUID as a varbinary.</li> <li>• Any other value for newid returns NULL.</li> </ul>

**Examples** **Example 1** Creates a table with varchar columns 32 bytes long, then uses newid with no arguments with the insert statement:

```
create table t (UUID varchar(32))
go
insert into t values (newid())
insert into t values (newid())
go
select * from t
```

```
UUID
-----
f81d4fae7dec11d0a76500a0c91e6bf6
7cd5b7769df75cefe040800208254639
```

**Example 2** Produces a GUID that includes dashes:

```
select newid(1)
-----
b59462af-a55b-469d-a79f-1d6c3c1e19e3
```

**Example 3** Returns a new GUID of type varbinary for every row that is returned from the query:

```
select newid(0x0) from sysobjects
```

**Example 4** Uses newid with the varbinary datatype:

```
sp_addtype binguid, "varbinary(16)"
create default binguid_dflt
```

```
as
newid(0x0)
sp_bindefault "binguid_dflt","binguid"
create table T1 (empname char(60), empid int, emp_guid
binguid)
insert T1 (empname, empid) values ("John Doe", 1)
insert T1 (empname, empid) values ("Jane Doe", 2)
```

Usage

- `newid` generates two values for the globally unique ID (GUID) based on arguments you pass to `newid`. The default argument generates GUIDs without dashes. By default `newid` returns new values for every filtered row.
- You can use `newid` in defaults, rules, and triggers, similar to other functions.
- Make sure the length of the varchar column is at least 32 bytes for the GUID format without dashes, and at least 36 bytes for the GUID format with dashes. The column length is truncated if it is not declared with these minimum required lengths. Truncation increases the probability of duplicate values.
- An argument of zero is equivalent to the default.
- Because GUIDs are globally unique, they can be transported across domains without generating duplicates.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute `newid`.

## next\_identity

Description	Retrieves the next identity value that is available for the next insert.
Syntax	<code>next_identity(table_name)</code>
Parameters	<i>table_name</i> identifies the table being used.
Examples	Updates the value of c2 to 10. The next available value is 11. <pre>select next_identity ("t1") t1 ----- 11</pre>
Usage	<ul style="list-style-type: none"><li>• <code>next_identity</code> returns the next value to be inserted by this task. In some cases, if multiple users are inserting values into the same table, the actual value reported as the next value to be inserted is different from the actual value inserted if another user performs an intermediate insert.</li><li>• <code>next_identity</code> returns a varchar character to support any precision of the identity column. If the table is a proxy table, a non-user table, or the table does not have identity property, NULL is returned.</li></ul>
Permissions	Only the table owner, System Administrator, or database administrator can issue this command.

## nullif

Description	Supports conditional SQL expressions; can be used anywhere a value expression can be used; alternative for a case expression.
Syntax	<code>nullif(<i>expression</i>, <i>expression</i>)</code>
Parameters	<code>nullif</code> compares the values of the two expressions. If the first expression equals the second expression, nullif returns NULL. If the first expression does not equal the second expression, nullif returns the first expression.  <i>expression</i> is a column name, a constant, a function, a subquery, or any combination of column names, constants, and functions connected by arithmetic or bitwise operators. For more information about expressions, see “Expressions” on page 275.

**Examples** **Example 1** Selects the *titles* and *type* from the *titles* table. If the book type is UNDECIDED, nullif returns a NULL value:

```
select title,  
       nullif(type, "UNDECIDED")  
from titles
```

**Example 2** This is an alternative way of writing Example 1:

```
select title,  
       case  
         when type = "UNDECIDED" then NULL  
         else type  
       end  
from titles
```

**Usage**

- nullif expression alternate for a case expression.
- nullif expression simplifies standard SQL expressions by allowing you to express a search condition as a simple comparison instead of using a when...then construct.
- You can use nullif expressions anywhere an expression can be used in SQL.
- At least one result of the case expression must return a non-null value. For example the following results in an error message:

```
select price, coalesce (NULL, NULL, NULL)  
from titles
```

All result expressions in a CASE expression must not be NULL.

- If your query produces a variety of datatypes, the datatype of a case expression result is determined by datatype hierarchy, as described in “Datatypes of mixed-mode expressions” on page 7. If you specify two datatypes that Adaptive Server cannot implicitly convert (for example, char and int), the query fails.

Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Anyone can execute nullif.
See also	<b>Commands</b> case, coalesce, select, if...else, where clause

## object\_id

Description	Returns the object ID of the specified object.
Syntax	<code>object_id(object_name)</code>
Parameters	<i>object_name</i> is the name of a database object, such as a table, view, procedure, trigger, default, or rule. The name can be fully qualified (that is, it can include the database and owner name). Enclose the <i>object_name</i> in quotes.
Examples	<b>Example 1</b> <pre>select object_id("titles") ----- 208003772</pre> <b>Example 2</b> <pre>select object_id("master..sysobjects") ----- 1</pre>
Usage	<ul style="list-style-type: none"><li>• <code>object_id</code>, a system function, returns the object's ID. Object IDs are stored in the <code>id</code> column of <code>sysobjects</code>.</li><li>• For general information about system functions, see "System functions" on page 68.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>object_id</code> .
See also	<b>Functions</b> <code>col_name</code> , <code>db_id</code> , <code>object_name</code> <b>System procedure</b> <code>sp_help</code>



## object\_name

Description	Returns the name of the object with the object ID you specify; can be up to 255 bytes in length.
Syntax	<code>object_name(object_id[, database_id])</code>
Parameters	<p><i>object_id</i> is the object ID of a database object, such as a table, view, procedure, trigger, default, or rule. Object IDs are stored in the <code>id</code> column of <code>sysobjects</code>.</p> <p><i>database_id</i> is the ID for a database if the object is not in the current database. Database IDs are stored in the <code>db_id</code> column of <code>sysdatabases</code>.</p>
Examples	<p><b>Example 1</b></p> <pre>select object_name(208003772) ----- titles</pre> <p><b>Example 2</b></p> <pre>select object_name(1, 1) ----- sysobjects</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>object_name</code>, a system function, returns the object's name.</li> <li>• For general information about system functions, see "System functions" on page 68.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>object_name</code> .
See also	<p><b>Functions</b> <code>col_name</code>, <code>db_id</code>, <code>object_id</code></p> <p><b>System procedure</b> <code>sp_help</code></p>

## pagesize

**Description** Returns the page size, in bytes, for the specified object.

**Syntax**

```
pagesize(object_name [, index_name])
pagesize(object_id [, db_id [, index_id]])
```

**Parameters**

*object\_name*  
is the object name of the page size of this function returns.

*index\_name*  
indicates the index name of the page size you want returned.

*object\_id*  
is the object ID of the page size this function returns.

*db\_id*  
is the database ID of the object.

*index\_id*  
is the index ID of the object you want returned.

**Examples** **Example 1** Selects the page size for the title\_id index in the current database.

```
select pagesize("title", "title_id")
```

**Example 2** The following returns the page size of the data layer for the object with *object\_id* 1234 and the database with a *db\_id* of 2 (the previous example defaults to the current database):

```
select pagesize(1234, 2, null)
select pagesize(1234, 2)
select pagesize(1234)
```

**Example 3** The following all default to the current database:

```
select pagesize(1234, null, 2)
select pagesize(1234)
```

**Example 4** Selects the page size for the titles table (object\_id 224000798) from the pubs2 database (db\_id 4):

```
select pagesize(224000798, 4)
```

**Example 5** Returns the page size for the nonclustered index's pages table mytable, residing in the current database:

```
pagesize(object_id('mytable'), NULL, 2)
```

**Example 6** Returns the page size for object titles\_clustindex from the current database:

	<pre>select pagesize("titles", "titles_clustindex")</pre>
Usage	<ul style="list-style-type: none"><li>• <code>pagesize</code> defaults to the data layer if you do not provide an index name or <i>index_id</i> (for example, <code>select pagesize("t1")</code>) if you use the word “null” as a parameter (for example, <code>select pagesize("t1", null)</code>).</li><li>• If the specified object is not an object requiring physical data storage for pages (for example, if you provide the name of a view), <code>pagesize</code> returns 0.</li><li>• If the specified object does not exist, <code>pagesize</code> returns NULL.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>pagesize</code> .

## partition\_id

**Description** Returns the partition ID of the specified data or index partition name.

**Syntax** `partition_id(table_name, partition_name [,index_name] )`

**Parameters** *table\_name*  
is the name for a table.

*partition\_name*  
is the partition name for a table partition or an index partition.

*index\_name*  
is the name of the index of interest.

**Examples** **Example 1** Returns the partition ID corresponding to the partition name testtable\_ptn1 and index id 0 (the base table). The testtable must exist in the current database:

```
select partition_id("testtable", "testtable_ptn1")
```

**Example 2** Returns the partition ID corresponding to the partition name testtable\_clust\_ptn1 for the index name clust\_index1. The testtable must exist in the current database:

```
select partition_id("testtable", "testtable_clust_ptn1", "clust_index1")
```

**Example 3** This is the same as the previous example, except that the user need not be in the same database as where the target table is located:

```
select partition_id("mydb.dbo.testtable", "testtable_clust_ptn1",  
"clust_index1")
```

**Usage** You must enclose *table\_name*, *partition\_name* and *index\_name* in quotes.

**See also** **Functions** data\_pages, object\_id, partition\_name, reserved\_pages, row\_count, used\_pages

## partition\_name

Description	The explicit name of a new partition, <code>partition_name</code> returns the partition name of the specified data or index partition id.
Syntax	<code>partition_name(indid, ptnid [, dbid])</code>
Parameters	<p><i>indid</i> is the index ID for the target partition.</p> <p><i>ptnid</i> is the ID of the target partition.</p> <p><i>dbid</i> is the database ID for the target partition. If you do not specify this parameter, the target partition is assumed to be in the current database.</p>
Examples	<p><b>Example 1</b> Returns the partition name for the given partition ID belonging to the base table (with an index ID of 0). The lookup is done in the current database because it does not specify a database ID:</p> <pre>select partition_name(0, 1111111111)</pre> <p><b>Example 2</b> Returns the partition name for the given partition ID belonging to the clustered index (index ID of 1 is specified) in the testdb database.</p> <pre>select partition_name(1, 1212121212, db_id("testdb"))</pre>
Usage	<ul style="list-style-type: none"> <li>If the search does not find the target partition, the return is NULL.</li> </ul>
See also	<b>Functions</b> <code>data_pages</code> , <code>object_id</code> , <code>partition_id</code> , <code>reserved_pages</code> , <code>row_count</code>

## patindex

Description	Returns the starting position of the first occurrence of a specified pattern.														
Syntax	<code>patindex("%pattern%", char_expr uchar_expr [, using {bytes   characters   chars} ] )</code>														
Parameters	<p><i>pattern</i> is a character expression of the char or varchar datatype that may include any of the pattern-match wildcard characters supported by Adaptive Server. The % wildcard character must precede and follow <i>pattern</i> (except when searching for first or last characters). For a description of the wildcard characters, see “Pattern matching with wildcard characters” on page 293.</p> <p><i>char_expr</i> is a character-type column name, variable, or constant expression of char, varchar, nchar, or nvarchar type.</p> <p><i>uchar_expr</i> is a character-type column name, variable, or constant expression of unichar, or univarchar type.</p> <p>using specifies a format for the starting position.</p> <p>bytes returns the offset in bytes.</p> <p>chars or characters returns the offset in characters (the default).</p>														
Examples	<p><b>Example 1</b> Selects the author ID and the starting character position of the word “circus” in the copy column:</p> <pre>select au_id, patindex("%circus%", copy) from blurbs</pre> <table><thead><tr><th>au_id</th><th></th></tr></thead><tbody><tr><td>486-29-1786</td><td>0</td></tr><tr><td>648-92-1872</td><td>0</td></tr><tr><td>998-72-3567</td><td>38</td></tr><tr><td>899-46-2035</td><td>31</td></tr><tr><td>672-71-3249</td><td>0</td></tr><tr><td>409-56-7008</td><td>0</td></tr></tbody></table> <p><b>Example 2</b></p> <pre>select au_id, patindex("%circus%", copy, using chars)</pre>	au_id		486-29-1786	0	648-92-1872	0	998-72-3567	38	899-46-2035	31	672-71-3249	0	409-56-7008	0
au_id															
486-29-1786	0														
648-92-1872	0														
998-72-3567	38														
899-46-2035	31														
672-71-3249	0														
409-56-7008	0														

```
from blurbs
```

**Example 3** Finds all the rows in sysobjects that start with “sys” with a fourth character that is “a”, “b”, “c”, or “d”:

```
select name
from sysobjects
where patindex("sys[a-d]%", name) > 0

name
-----
sysalternates
sysattributes
syscharsets
syscolumnms
syscomments
sysconfigures
sysconstraints
syscurconfigs
sysdatabases
sysdepends
sysdevices
```

#### Usage

- patindex, a string function, returns an integer representing the starting position of the first occurrence of *pattern* in the specified character expression, or a 0 if *pattern* is not found.
- You can use patindex on all character data, including text and image data.
- For unichar, univarchar, and unitext, patindex returns the offset in Unicode characters. The pattern string is implicitly converted to UTF-16 before comparison, and the comparison is based on the default unicode sort order configuration. For example, this is what is returned if a unitext column contains row value U+0041U+0042U+d800U+dc00U+0043:

```
select patindex("%C%", ut) from unitable
-----
4
```

- By default, patindex returns the offset in characters; to return the offset in bytes (multibyte character strings), specify using bytes.
- Include percent signs before and after *pattern*. To look for *pattern* as the first characters in a column, omit the preceding %. To look for *pattern* as the last characters in a column, omit the trailing %.
- If *char\_expr* or *uchar\_expr* is NULL, patindex returns 0.

- If you give a varchar expression as one parameter and a unichar expression as the other, the varchar expression is implicitly converted to unichar (with possible truncation).
- For general information about string functions, see “String functions” on page 67.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute `patindex`.

See also

**Functions** `charindex`, `substring`



## pi

Description	Returns the constant value 3.1415926535897936.
Syntax	pi()
Parameters	None
Examples	<pre>select pi() ----- 3.141593</pre>
Usage	<ul style="list-style-type: none"><li>• pi, a mathematical function, returns the constant value of 3.1415926535897931.</li><li>• For general information about mathematical functions, see “Mathematical functions” on page 65.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute pi.
See also	<b>Functions</b> degrees, radians

## power

Description	Returns the value that results from raising the specified number to a given power.
Syntax	<code>power(value, power)</code>
Parameters	<i>value</i> is a numeric value.  <i>power</i> is an exact numeric, approximate numeric, or money value.
Examples	<pre>select power(2, 3) -----             8</pre>
Usage	<ul style="list-style-type: none"><li>power, a mathematical function, returns the value of <i>value</i> raised to the power <i>power</i>. Results are of the same type as <i>value</i>.  In expressions of type numeric or decimal, this function returns precision:38, scale 18.</li><li>For general information about mathematical functions, see “Mathematical functions” on page 65.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute power.
See also	<b>Functions</b> exp, log, log10

## proc\_role

Description	<p>Returns information about whether the user has been granted the specified role.</p> <hr/> <p><b>Note</b> Sybase supports—and recommends—that you use <code>has_role</code> instead of <code>proc_role</code>. You need not, however, convert your existing uses of <code>proc_role</code> to <code>has_role</code>.</p> <hr/>
Syntax	<code>proc_role ("role_name")</code>
Parameters	<p><i>role_name</i> is the name of a system or user-defined role.</p>
Examples	<p><b>Example 1</b> Creates a procedure to check if the user is a System Administrator:</p> <pre>create procedure sa_check as if (proc_role("sa_role") &gt; 0) begin     print "You are a System Administrator."     return(1) end</pre> <p><b>Example 2</b> Checks that the user has been granted the System Security Officer role:</p> <pre>select proc_role("sso_role")</pre> <p><b>Example 3</b> Checks that the user has been granted the Operator role:</p> <pre>select proc_role("oper_role")</pre>
Usage	<ul style="list-style-type: none"> <li>• Using <code>proc_role</code> with a procedure that starts with “<code>sp_</code>” returns an error.</li> <li>• <code>proc_role</code>, a system function, checks whether an invoking user has been granted, and has activated, the specified role.</li> <li>• <code>proc_role</code> returns 0 if the user has:             <ul style="list-style-type: none"> <li>• Not been granted the specified role</li> <li>• Not been granted a role which contains the specified role</li> <li>• Been granted, but has not activated, the specified role</li> </ul> </li> <li>• <code>proc_role</code> returns 1 if the invoking user has been granted, and has activated, the specified role.</li> <li>• <code>proc_role</code> returns 2 if the invoking user has a currently active role, which contains the specified role.</li> </ul>

- For general information about system functions, see “System functions” on page 68.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute `proc_role`.

See also

**Commands** `alter role`, `create role`, `drop role`, `grant`, `set`, `revoke`

**Functions** `mut_excl_roles`, `role_contain`, `role_id`, `role_name`, `show_role`

## radians

Description	Returns the size, in radians, of an angle with the specified number of degrees.
Syntax	<code>radians(<i>numeric</i>)</code>
Parameters	<p><i>numeric</i></p> <p>is any exact numeric (numeric, dec, decimal, tinyint, smallint, or int), approximate numeric (float, real, or double precision), or money column, variable, constant expression, or a combination of these.</p>
Examples	<pre>select radians(2578) -----                 44</pre>
Usage	<ul style="list-style-type: none"> <li>radians, a mathematical function, converts degrees to radians. Results are of the same type as <i>numeric</i>.</li> </ul> <p>To express numeric or decimal datatypes, this function returns precision: 38, scale 18.</p> <p>When money datatypes are used, internal conversion to float may cause loss of precision.</p> <ul style="list-style-type: none"> <li>For general information about mathematical functions, see “Mathematical functions” on page 65.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute radians.
See also	<b>Function</b> degrees

## rand

Description	Returns a random value between 0 and 1, which is generated using the specified seed value.
Syntax	<code>rand([integer])</code>
Parameters	<i>integer</i> is any integer (tinyint, smallint, or int) column name, variable, constant expression, or a combination of these.
Examples	<p><b>Example 1</b></p> <pre>select rand() ----- 0.395740</pre> <p><b>Example 2</b></p> <pre>declare @seed int select @seed=100 select rand(@seed) ----- 0.000783</pre>
Usage	<ul style="list-style-type: none"><li>• rand, a mathematical function, returns a random float value between 0 and 1, using the optional integer as a seed value.</li><li>• The rand function uses the output of a 32-bit pseudorandom integer generator. The integer is divided by the maximum 32-bit integer to give a double value between 0.0 and 1.0. The rand function is seeded randomly at server start-up, so getting the same sequence of random numbers is unlikely, unless the user first initializes this function with a constant seed value. The rand function is a global resource. Multiple users calling the rand function progress along a single stream of pseudorandom values. If a repeatable series of random numbers is needed, the user must assure that the function is seeded with the same value initially and that no other user calls rand while the repeatable sequence is desired.</li><li>• For general information about mathematical functions, see “Mathematical functions” on page 65.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute rand.
See also	<b>Datatypes</b> Approximate numeric datatypes

## replicate

Description	Returns a string consisting of the specified expression repeated a given number of times.
Syntax	<code>replicate (char_expr   uchar_expr, integer_expr)</code>
Parameters	<p><i>char_expr</i> is a character-type column name, variable, or constant expression of char, varchar, nchar, or nvarchar type.</p> <p><i>uchar_expr</i> is a character-type column name, variable, or constant expression of unichar or univarchar type.</p> <p><i>integer_expr</i> is any integer (tinyint, smallint, or int) column name, variable, or constant expression.</p>
Examples	<pre>select replicate("abcd", 3) ----- abcdabcdabcd</pre>
Usage	<ul style="list-style-type: none"> <li>• replicate, a string function, returns a string with the same datatype as <i>char_expr</i> or <i>uchar_expr</i> containing the same expression repeated the specified number of times or as many times as fits into 16K, whichever is less.</li> <li>• If <i>char_expr</i> or <i>uchar_expr</i> is NULL, returns a single NULL.</li> <li>• For general information about string functions, see “String functions” on page 67.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute replicate.
See also	<b>Function</b> stuff

## reserved\_pages

Description	<p>Reports the number of pages reserved to a table, index or a specific partition. The result includes pages used for internal structures.</p> <p>This function replaces the old reserved_pgs function used in Adaptive Server versions earlier than 15.0.</p>
Syntax	<code>reserved_pages(<i>dbid</i>, <i>object_id</i> [, <i>indid</i> [, <i>ptnid</i>]])</code>
Parameters	<p><i>dbid</i> is the database ID of the database where the target object resides.</p> <p><i>object_id</i> is an object ID for a table.</p> <p><i>indid</i> is the index ID of target index.</p> <p><i>ptnid</i> is the partition ID of target partition.</p>
Examples	<p><b>Example 1</b> Returns the number of pages reserved by the object with a object ID of 31000114 in the specified database (including any indexes):</p> <pre>select reserved_pages(5, 31000114)</pre> <p><b>Example 2</b> Returns the number of pages reserved by the object in the data layer, regardless of whether or not a clustered index exists:</p> <pre>select reserved_pages(5, 31000114, 0)</pre> <p><b>Example 3</b> Returns the number of pages reserved by the object in the index layer for a clustered index. This does not include the pages used by the data layer:</p> <pre>select reserved_pages(5, 31000114, 1)</pre> <p><b>Example 4</b> Returns the number of pages reserved by the object in the data layer of the specific partition, which in this case is 2323242432:</p> <pre>select reserved_pages(5, 31000114, 0, 2323242432)</pre>
Usage	<p>In the case of an apl table, if a clustered index exists on the table, then passing in an in did of 0 will report the reserved data pages, and passing an indid of 1 will report the reserved index pages. All erroneous conditions will result in a value of zero being returned.</p>
Standards	<p>ANSI SQL – Compliance level: Transact-SQL extension.</p>
Permissions	<p>Any user can execute reserved_pgs.</p>



See also

**Command** update statistics

**Function** data\_pages, reserved\_pages, row\_count, used\_pages

## reverse

Description	Returns the specified string with characters listed in reverse order.
Syntax	<code>reverse(<i>expression</i>   <i>uchar_expr</i>)</code>
Parameters	<p><i>expression</i></p> <p>is a character or binary-type column name, variable, or constant expression of char, varchar, nchar, nvarchar, binary, or varbinary type.</p> <p><i>uchar_expr</i></p> <p>is a character or binary-type column name, variable, or constant expression of unichar or univarchar type.</p>
Examples	<p><b>Example 1</b></p> <pre>select reverse("abcd") ----- dcba</pre> <p><b>Example 2</b></p> <pre>select reverse(0x12345000) ----- 0x00503412</pre>
Usage	<ul style="list-style-type: none"><li>• <code>reverse</code>, a string function, returns the reverse of <i>expression</i>.</li><li>• If <i>expression</i> is NULL, <code>reverse</code> returns NULL.</li><li>• Surrogate pairs are treated as indivisible and are not reversed.</li><li>• For general information about string functions, see “String functions” on page 67.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>reverse</code> .
See also	<b>Functions</b> lower, upper

## right

Description	The rightmost part of the expression with the specified number of characters.
Syntax	<code>right(expression, integer_expr)</code>
Parameters	<p><i>expression</i> is a character or binary-type column name, variable, or constant expression of char, varchar, nchar, unichar, nvarchar, univarchar, binary, or varbinary type.</p> <p><i>integer_expr</i> is any integer (tinyint, smallint, or int) column name, variable, or constant expression.</p>
Examples	<p><b>Example 1</b></p> <pre>select right("abcde", 3) --- cde</pre> <p><b>Example 2</b></p> <pre>select right("abcde", 2) -- de</pre> <p><b>Example 3</b></p> <pre>select right("abcde", 6) ----- abcde</pre> <p><b>Example 4</b></p> <pre>select right(0x12345000, 3) ----- 0x345000</pre> <p><b>Example 5</b></p> <pre>select right(0x12345000, 2) ----- 0x5000</pre> <p><b>Example 6</b></p> <pre>select right(0x12345000, 6) ----- 0x12345000</pre>

Usage	<ul style="list-style-type: none"><li>• <i>right</i>, a string function, returns the specified number of characters from the rightmost part of the character or binary expression.</li><li>• If the specified rightmost part begins with the second surrogate of a pair (the low surrogate), the return value starts with the next full character. Therefore, one less character is returned.</li><li>• The return value has the same datatype as the character or binary expression.</li><li>• If <i>expression</i> is NULL, <i>right</i> returns NULL.</li><li>• For general information about string functions, see “String functions” on page 67.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension
Permissions	Any user can execute <i>right</i> .
See also	<b>Functions</b> <i>rtrim</i> , <i>substring</i>

## rm\_appcontext

Description	Removes a specific application context, or all application contexts. <code>rm_appcontext</code> is a function provided by the Application Context Facility (ACF).
Syntax	<code>rm_appcontext ("context_name", "attribute_name")</code>
Parameters	<p><i>context_name</i> is a row specifying an application context name. It is saved as datatype <code>char(30)</code>.</p> <p><i>attribute_name</i> is a row specifying an application context attribute name. It is saved as datatype <code>char(30)</code>.</p>
Examples	<p><b>Example 1</b> Removes an application context by specifying some or all attributes:</p> <pre>select rm_appcontext ("CONTEXT1", "*") ----- 0  select rm_appcontext ("*", "*") ----- 0  select rm_appcontext ("NON_EXISTING_CTX", "ATTR") ----- -1</pre> <p><b>Example 2</b> Shows the result when a user without appropriate permissions attempts to remove an application context:</p> <pre>select rm_appcontext ("CONTEXT1", "ATTR2") ----- -1</pre>
Usage	<ul style="list-style-type: none"> <li>• This function always returns 0 for success.</li> <li>• All the arguments for this function are required.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Permissions depend on the user profile and the application profile, which are stored by ACF.
See also	For more information on the ACF see “Row-level access control” in Chapter 11, “Managing User Permissions” of the <i>System Administration Guide</i> .
	<b>Functions</b> <code>get_appcontext</code> , <code>list_appcontext</code> , <code>set_appcontext</code>

## role\_contain

**Description** Returns 1 if *role2* contains *role1*.

**Syntax** `role_contain("role1", "role2")`

**Parameters**

*role1*  
is the name of a system or user-defined role.

*role2*  
is the name of another system or user-defined role.

**Examples**

**Example 1**

```
select role_contain("intern_role", "doctor_role")
-----
1
```

**Example 2**

```
select role_contain("specialist_role", "intern_role")
-----
0
```

**Usage**

- `role_contain`, a system function, returns 1 if *role1* is contained by *role2*.
- For more information about system functions, see “System functions” on page 68.

**Standards** ANSI SQL – Compliance level: Transact-SQL extension.

**Permissions** Any user can execute `role_contain`.

**See also** **Documents** For more information about contained roles and role hierarchies, see the *System Administration Guide*.

**Functions** `mut_excl_roles`, `proc_role`, `role_id`, `role_name`

**Commands** `alter role`

**System procedures** `sp_activeroles`, `sp_displayroles`, `sp_role`

## role\_id

Description	Returns the system role ID of the name you specify.
Syntax	<code>role_id("role_name")</code>
Parameters	<p><i>role_name</i></p> <p>is the name of a system or user-defined role. Role names and role IDs are stored in the <code>sysrvroles</code> system table.</p>
Examples	<p><b>Example 1</b> Returns the system role ID of <code>sa_role</code>:</p> <pre>select role_id("sa_role") ----- 0</pre> <p><b>Example 2</b> Returns the system role ID of the “<code>intern_role</code>”:</p> <pre>select role_id("intern_role") ----- 6</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>role_id</code>, a system function, returns the system role ID (<code>srld</code>). System role IDs are stored in the <code>srld</code> column of the <code>sysrvroles</code> system table.</li> <li>• If the <i>role_name</i> is not a valid role in the system, Adaptive Server returns NULL.</li> <li>• For more information about system functions, see “System functions” on page 68.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>role_id</code> .
See also	<p><b>Documents</b> For more information about roles, see the <i>System Administration Guide</i>.</p> <p><b>Functions</b> <code>mut_excl_roles</code>, <code>proc_role</code>, <code>role_contain</code>, <code>role_name</code></p>

## **role\_name**

Description	Returns the name of a system role ID you specify.
Syntax	<code>role_name(role_id)</code>
Parameters	<i>role_id</i> is the system role ID (srid) of the role. Role names are stored in sysssvroles.
Examples	<pre>select role_name(01) ----- sso_role</pre>
Usage	<ul style="list-style-type: none"><li>• <code>role_name</code>, a system function, returns the role name.</li><li>• For more information about system functions, see “System functions” on page 68.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension
Permissions	Any user can execute <code>role_name</code> .
See also	<b>Functions</b> <code>mut_excl_roles</code> , <code>proc_role</code> , <code>role_contain</code> , <code>role_id</code>



# round

**Description** Returns the value of the specified number, rounded to a specified number of decimal places.

**Syntax** `round(number, decimal_places)`

**Parameters** *number*  
is any exact numeric (numeric, dec, decimal, tinyint, smallint, int, or bigint), approximate numeric (float, real, or double precision), or money column, variable, constant expression, or a combination of these.

*decimal\_places*  
is the number of decimal places to round to.

## Examples

### Example 1

```
select round(123.4545, 2)
-----
123.4500
```

### Example 2

```
select round(123.45, -2)
-----
100.00
```

### Example 3

```
select round(1.2345E2, 2)
-----
123.450000
```

### Example 4

```
select round(1.2345E2, -2)
-----
100.000000
```

## Usage

- `round`, a mathematical function, rounds the *number* so that it has *decimal\_places* significant digits.
- A positive value for *decimal\_places* determines the number of significant digits to the right of the decimal point; a negative value for *decimal\_places* determines the number of significant digits to the left of the decimal point.
- Results are of the same type as *number* and, for numeric and decimal expressions, have an internal precision equal to the precision of the first argument plus 1 and a scale equal to that of *number*.

- `round` always returns a value. If *decimal\_places* is negative and exceeds the number of significant digits specified for *number*, Adaptive Server returns 0. (This is expressed in the form 0.00, where the number of zeros to the right of the decimal point is equal to the scale of numeric.) For example, the following returns a value of 0.00:

```
select round(55.55, -3)
```

- For general information about mathematical functions, see “Mathematical functions” on page 65.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute `round`.

See also

**Functions** `abs`, `ceiling`, `floor`, `sign`, `str`

## row\_count

Description	Returns an estimate of the number of rows in the specified table.
Syntax	<code>row_count(<i>dbid</i>, <i>object_id</i> [,<i>ptnid</i>])</code>
Parameters	<i>dbid</i> the database ID where target object resides  <i>object_id</i> object ID of table  <i>ptnid</i> partition ID of interest
Examples	<b>Example 1</b> Returns an estimate of the number of rows in the given object: <pre>select row_count(5, 31000114)</pre> <b>Example 2</b> Returns an estimate of the number of rows in the specified partition (with partition ID of 2323242432) of the object with object ID of 31000114: <pre>select row_count(5, 31000114, 2323242432)</pre>
Usage	All erroneous conditions will return in a value of zero being returned.
Standards	ANSI SQL – Compliance level: Transact-SQL extension
Permissions	Any user can execute row_count.
See also	<b>Functions</b> reserved_pages, used_pages

## rtrim

Description	Returns the specified expression, trimmed of trailing blanks.
Syntax	<code>rtrim(<i>char_expr</i>   <i>uchar_expr</i>)</code>
Parameters	<p><i>char_expr</i> is a character-type column name, variable, or constant expression of char, varchar, nchar, or nvarchar type.</p> <p><i>uchar_expr</i> is a character-type column name, variable, or constant expression of unichar or univarchar type.</p>
Examples	<pre>select rtrim("abcd  ") ----- abcd</pre>
Usage	<ul style="list-style-type: none"><li>• rtrim, a string function, removes trailing blanks.</li><li>• For Unicode, a blank is defined as the Unicode value U+0020.</li><li>• If <i>char_expr</i> or <i>uchar_expr</i> is NULL, returns NULL.</li><li>• Only values equivalent to the space character in the current character set are removed.</li><li>• For general information about string functions, see “String functions” on page 67.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute rtrim.
See also	<b>Function</b> ltrim

## set\_appcontext

Description	Sets an application context name, attribute name, and attribute value for a user session, defined by the attributes of a specified application. <code>set_appcontext</code> is a built-in function that the Application Context Facility (ACF) provides.
Syntax	<code>set_appcontext ("context_name", "attribute_name", "attribute_value")</code>
Parameters	<p><i>context_name</i> is a row that specifies an application context name. It is saved as the datatype <code>char(30)</code>.</p> <p><i>attribute_name</i> is a row that specifies an application context attribute name. It is saved as the datatype <code>char(30)</code>.</p> <p><i>attribute_value</i> is a row that specifies and application attribute value. It is saved as the datatype <code>char(30)</code>.</p>
Examples	<p><b>Example 1</b> Creates an application context called <code>CONTEXT1</code>, with an attribute <code>ATTR1</code> that has the value <code>VALUE1</code>.</p>

```
select set_appcontext ("CONTEXT1", "ATTR1", "VALUE1")
-----
0
```

Attempting to override the existing application context created causes the following:

```
select set_appcontext ("CONTEXT1", "ATTR1", "VALUE1")
-----
-1
```

**Example 2** Shows `set_appcontext` including a datatype conversion in the value.

```
declare @numericvarchar varchar(25)
select @numericvar = "20"
select set_appcontext ("CONTEXT1", "ATTR2",
convert(char(20), @numericvar))
-----
0
```

**Example 3** Shows the result when a user without appropriate permissions attempts to set the application context.

```
select set_appcontext ("CONTEXT1", "ATTR2", "VALUE1")
-----
```

-1

Usage	<ul style="list-style-type: none"><li>• set_appcontext returns 0 for success and -1 for failure.</li><li>• If you set values that already exist in the current session, set_appcontext returns -1.</li><li>• This function cannot override the values of an existing application context. To assign new values to a context, remove the context and re-create it using new values.</li><li>• set_appcontext saves attributes as char datatypes. If you are creating an access rule that must compare the attribute value to another datatype, the rule should convert the char data to the appropriate datatype.</li><li>• All the arguments for this function are required.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Permissions depend on the user profile and the application profile, stored by ACF.
See also	For more information on the ACF see “Row-level access control” in Chapter 11, “Managing User Permissions” of the <i>System Administration Guide</i> .
	<b>Functions</b> get_appcontext, list_appcontext, rm_appcontext

## show\_role

Description	Shows the login's currently active system-defined roles.
Syntax	show_role()
Parameters	None.
Examples	<p><b>Example 1</b></p> <pre>select show_role() sa_role sso_role oper_role replication_role</pre> <p><b>Example 2</b></p> <pre>if charindex("sa_role", show_role()) &gt;0 begin     print "You have sa_role" end</pre>
Usage	<ul style="list-style-type: none"> <li>• show_role, a system function, returns the login's current active system-defined roles, if any (sa_role, sso_role, oper_role, or replication_role). If the login has no roles, show_role returns NULL.</li> <li>• When a Database Owner invokes show_role after using setuser, show_role displays the active roles of the Database Owner, not the user impersonated with setuser.</li> <li>• For general information about system functions, see "System functions" on page 68.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute show_role.
See also	<p><b>Commands</b> alter role, create role, drop role, grant, set, revoke</p> <p><b>Functions</b> proc_role, role_contain</p> <p><b>System procedures</b> sp_active roles, sp_displayroles, sp_role</p>

## show\_sec\_services

Description	Lists the security services that are active for the session.
Syntax	show_sec_services()
Parameters	None.
Examples	Shows that the user's current session is encrypting data and performing replay detection checks: <pre>select show_sec_services() encryption, replay_detection</pre>
Usage	<ul style="list-style-type: none"><li>• Use show_sec_services to list the security services that are active during the session.</li><li>• If no security services are active, show_sec_services returns NULL.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute show_sec_services.
See also	<b>Functions</b> is_sec_service_on



# sign

Description	Returns the sign (1 for positive, 0, or -1 for negative) of the specified value.
Syntax	<code>sign(<i>numeric</i>)</code>
Parameters	<p><i>numeric</i></p> <p>is any exact numeric (numeric, dec, decimal, tinyint, smallint, int, or bigint), approximate numeric (float, real, or double precision), or money column, variable, constant expression, or a combination of these.</p>
Examples	<p><b>Example 1</b></p> <pre>select sign(-123) ----- -1</pre> <p><b>Example 2</b></p> <pre>select sign(0) ----- 0</pre> <p><b>Example 3</b></p> <pre>select sign(123) ----- 1</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>sign</code>, a mathematical function, returns the positive (1), zero (0), or negative (-1).</li> <li>• Results are of the same type, and have the same precision and scale, as the numeric expression.</li> <li>• For general information about mathematical functions, see “Mathematical functions” on page 65.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>sign</code> .
See also	<b>Functions</b> <code>abs</code> , <code>ceiling</code> , <code>floor</code> , <code>round</code>

## sin

Description	Returns the sine of the specified angle (in radians).
Syntax	<code>sin(<i>approx_numeric</i>)</code>
Parameters	<i>approx_numeric</i> is any approximate numeric (float, real, or double precision) column name, variable, or constant expression.
Examples	<pre>select sin(45) -----           0.850904</pre>
Usage	<ul style="list-style-type: none"><li>• <code>sin</code>, a mathematical function, returns the sine of the specified angle (measured in radians).</li><li>• For general information about mathematical functions, see “Mathematical functions” on page 65.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>sin</code> .
See also	<b>Functions</b> <code>cos</code> , degrees, radians

## sortkey

Description	Generates values that can be used to order results based on collation behavior, which allows you to work with character collation behaviors beyond the default set of Latin character-based dictionary sort orders and case- or accent-sensitivity.
Syntax	<code>sortkey (char_expression   uchar_expression) [, {collation_name   collation_ID}]</code>
Parameters	<p><i>char_expression</i> is a character-type column name, variable, or constant expression of char, varchar, nchar, or nvarchar type.</p> <p><i>uchar_expression</i> is a character-type column name, variable, or constant expression of unichar or univarchar type.</p> <p><i>collation_name</i> is a quoted string or a character variable that specifies the collation to use. Table 2-10 on page 226 shows the valid values.</p> <p><i>collation_ID</i> is an integer constant or a variable that specifies the collation to use. Table 2-10 on page 226 shows the valid values.</p>
Examples	<p><b>Example 1</b> Shows sorting by European language dictionary order:</p> <pre>select * from cust_table where cust_name like "TI%" order by (sortkey(cust_name, "dict"))</pre> <p><b>Example 2</b> Shows sorting by simplified Chinese phonetic order:</p> <pre>select *from cust_table where cust name like "TI%" order by (sortkey(cust-name, "gbpinyin"))</pre> <p><b>Example 3</b> Shows sorting by European language dictionary order using the in-line option:</p> <pre>select *from cust_table where cust_name like "TI%" order by cust_french_sort</pre> <p><b>Example 4</b> Shows sorting by Simplified Chinese phonetic order using preexisting keys:</p> <pre>select * from cust_table where cust_name like "TI%" order by cust_chinese_sort.</pre>

Usage

- sortkey, a system function, generates values that can be used to order results based on collation behavior. This allows you to work with character collation behaviors beyond the default set of Latin-character-based dictionary sort orders and case- or accent-sensitivity. The return value is a varbinary datatype value that contains coded collation information for the input string that is returned from the sortkey function.

For example, you can store the values returned by sortkey in a column with the source character string. To retrieve the character data in the desired order, include in the select statement an order by clause on the columns that contain the results of running sortkey.

sortkey guarantees that the values it returns for a given set of collation criteria work for the binary comparisons that are performed on varbinary datatypes.

- sortkey can generate up to sixbytes of collation information for each input character. Therefore, the result from using sortkey may exceed the length limit of the varbinary datatype. If this happens, the result is truncated to fit. Since this limit is dependent on the logical page size of your server, truncation removes result bytes for each input character until the result string is less than the following for DOL and APL tables:

**Table 2-9: Maximum row and column length—APL and DOL tables**

Locking scheme	Page size	Maximum row length	Maximum column length
APL tables	2K (2048 bytes)	1962	1960 bytes
	4K (4096 bytes)	4010	4008 bytes
	8K (8192 bytes)	8106	8104 bytes
	16K (16384 bytes)	16298	16296 bytes
DOL tables	2K (2048 bytes)	1964	1958 bytes
	4K (4096 bytes)	4012	4006 bytes
	8K (8192 bytes)	8108	8102 bytes
	16K (16384 bytes)	16300	16294 bytes if table does not include any variable length columns
	16K (16384 bytes)	16300 (subject to a max start offset of varlen = 8191)	8191-6-2 = 8183 bytes if table includes at least one variable length column.*

\* This size includes six bytes for the row overhead and two bytes for the row length field.

If this occurs, Adaptive Server issues a warning message, but the query or transaction that contained the sortkey function continues to run.

- *char\_expression* or *uchar\_expression* must be composed of characters that are encoded in the server's default character set.
- *char\_expression* or *uchar\_expression* can be an empty string. If it is an empty string, sortkey returns a zero-length varbinary value, and stores a blank for the empty string.

An empty string has a different collation value than a NULL string from a database column.

- If *char\_expression* or *uchar\_expression* is NULL, sortkey returns a null value.
- If a unicode expression has no specified sort order, the unicode default sort order is used.
- If you do not specify a value for *collation\_name* or *collation\_ID*, sortkey assumes binary collation.
- The binary values generated from the sortkey function can change from one major version to another major version of Adaptive Server, such as version 12.0 to 12.5, version 12.9.2 to 12.0, and so on. If you are upgrading to the current version of Adaptive Server, regenerate keys and repopulate the shadow columns before any binary comparison takes place.

---

**Note** Upgrades from version 12.5 to 12.5.0.1 do not require this step, and Adaptive Server does not generate any errors or warning messages if you do not regenerate the keys. Although a query involving the shadow columns should work fine, the comparison result may differ from the pre-upgrade server.

---

### Collation tables

There are two types of collation tables you can use to perform multilingual sorting:

- 1 A "built-in" collation table created by the sortkey function. This function exists in versions of Adaptive Server later than 11.5.1. You can use either the collation name or the collation ID to specify a built-in table.
- 2 An external collation table that uses the Unilib library sorting functions. You must use the collation name to specify an external table. These files are located in *\$SYBASE/collate/unicode*.

Both of these methods work equally well, but a “built-in” table is tied to a Adaptive Server database, while an external table is not. If you use an Adaptive Server database, a built-in table provides the best performance. Both methods can handle any mix of English, European, and Asian languages.

There are two ways to use sortkey:

- 1 In-line – this uses sortkey as part of the order by clause and is useful for retrofitting an existing application and minimizing the changes. However, this method generates sort keys on-the-fly, and therefore does not provide optimum performance on large data sets of more than 1000 records.
- 2 Pre-existing keys – this method calls sortkey whenever a new record requiring multilingual sorting is added to the table, such as a new customer name. Shadow columns (binary or varbinary type) must be set up in the database, preferably in the same table, one for each desired sort order such as French, Chinese, and so on. When a query requires output to be sorted, the order by clause uses one of the shadow columns. This method produces the best performance since keys are already generated and stored, and are quickly compared only on the basis of their binary values.

You can view a list of available collation rules. Print the list by executing either `sp_helpsort`, or by querying and selecting the name, id, and description from `syscharsets` (type is between 2003 and 2999).

- Table 2-10 lists the valid values for *collation\_name* and *collation\_ID*.

**Table 2-10: Collation names and IDs**

Description	Collation name	Collation ID
Default Unicode multilingual	default	20
Thai dictionary order	thaidict	21
ISO14651 standard	iso14651	22
UTF-16 ordering – matches UTF-8 binary ordering	utf8bin	24
CP 850 Alternative – no accent	altnoacc	39
CP 850 Alternative – lowercase first	altdict	45
CP 850 Western European – no case preference	altnocsp	46
CP 850 Scandinavian – dictionary ordering	scandict	47
CP 850 Scandinavian – case-insensitive with preference	scannocp	48
GB Pinyin	gbpinyin	n/a
Binary sort	binary	50
Latin-1 English, French, German dictionary	dict	51
Latin-1 English, French, German no case	nocase	52

Description	Collation name	Collation ID
Latin-1 English, French, German no case, preference	nocasep	53
Latin-1 English, French, German no accent	noaccent	54
Latin-1 Spanish dictionary	espdict	55
Latin-1 Spanish no case	espnocs	56
Latin-1 Spanish no accent	espnoac	57
ISO 8859-5 Russian dictionary	rusdict	58
ISO 8859-5 Russian no case	rusnocs	59
ISO 8859-5 Cyrillic dictionary	cyrdict	63
ISO 8859-5 Cyrillic no case	cyrnocs	64
ISO 8859-7 Greek dictionary	elldict	65
ISO 8859-2 Hungarian dictionary	hundict	69
ISO 8859-2 Hungarian no accents	hunnoac	70
ISO 8859-2 Hungarian no case	hunnocs	71
ISO 8859-9 Turkish dictionary	turdict	72
ISO 8859-9 Turkish no accents	turknoac	73
ISO 8859-9 Turkish no case	turknocs	74
CP932 binary ordering	cp932bin	129
Chinese phonetic ordering	dynix	130
GB2312 binary ordering	gb2312bn	137
Common Cyrillic dictionary	cyrdict	140
Turkish dictionary	turdict	155
EUCKSC binary ordering	euckscbn	161
Chinese phonetic ordering	gbpinyin	163
Russian dictionary ordering	rusdict	165
SJIS binary ordering	sjisbin	179
EUCJIS binary ordering	eucjisbn	192
BIG5 binary ordering	big5bin	194
Shift-JIS binary order	sjisbin	259

Standards ANSI SQL – Compliance level: Transact-SQL extension.

Permissions Any user can execute sortkey.

See also **Function** compare

## soundex

Description	Returns a four-character code representing the way an expression sounds.
Syntax	<code>soundex(char_expr   uchar_expr)</code>
Parameters	<p><i>char_expr</i> is a character-type column name, variable, or constant expression of char, varchar, nchar, or nvarchar type.</p> <p><i>uchar_expr</i> is a character-type column name, variable, or constant expression of unichar or univarchar type.</p>
Examples	<pre>select soundex ("smith"), soundex ("smythe") ----- S530  S530</pre>
Usage	<ul style="list-style-type: none"><li>• <code>soundex</code>, a string function, returns a four-character soundex code for character strings that are composed of a contiguous sequence of valid single- or double-byte roman letters.</li><li>• The <code>soundex</code> function converts an alphabetic string to a four-digit code for use in locating similar-sounding words or names. All vowels are ignored unless they constitute the first letter of the string.</li><li>• If <i>char_expr</i> or <i>uchar_expr</i> is NULL, returns NULL.</li><li>• For general information about string functions, see “String functions” on page 67.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>soundex</code> .
See also	<b>Function</b> difference



## space

Description	Returns a string consisting of the specified number of single-byte spaces.
Syntax	<code>space(integer_expr)</code>
Parameters	<i>integer_expr</i> is any integer (tinyint, smallint, or int) column name, variable, or constant expression.
Examples	<pre>select "aaa", space(4), "bbb"       ---  ----  ---       aaa      bbb</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>space</code>, a string function, returns a string with the indicated number of single-byte spaces.</li> <li>• For general information about string functions, see “String functions” on page 67.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>space</code> .
See also	<b>Functions</b> <code>isnull</code> , <code>rtrim</code>

## square

**Description** Returns the square of a specified value expressed as a float.

**Syntax** `square(numeric_expression)`

**Parameters** *numeric\_expression*  
is a numeric expression of type float.

**Examples** **Example 1** Returns the square from an integer column:

```
select square(total_sales) from titles
-----
16769025.00000
15023376.00000
350513284.00000
...
16769025.00000
(18 row(s) affected)
```

**Example 2** Returns the square from a money column:

```
select square(price) from titles
-----
399.600100
142.802500
8.940100
NULL
...
224.700100
(18 row(s) affected)
```

**Usage** This function is the equivalent of `power(numeric_expression,2)`, but it returns type float rather than int.

**Standards** ANSI SQL – Compliance level: Transact-SQL extension.

**Permissions** Any user can execute square.

**See also** **Function** power

**Datatypes** exact\_numeric, approximate\_numeric, money, float

## sqrt

Description	Returns the square root of the specified number.
Syntax	<code>sqrt(<i>approx_numeric</i>)</code>
Parameters	<i>approx_numeric</i> is any approximate numeric (float, real, or double precision) column name, variable, or constant expression that evaluates to a positive number.
Examples	<pre>select sqrt(4)         2.000000</pre>
Usage	<ul style="list-style-type: none"><li>• <code>sqrt</code>, a mathematical function, returns the square root of the specified value.</li><li>• If you attempt to select the square root of a negative number, Adaptive Server returns the following error message: <pre>Domain error occurred.</pre></li><li>• For general information about mathematical functions, see “Mathematical functions” on page 65.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>sqrt</code> .
See also	<b>Function</b> <code>power</code>

## str

Description	Returns the character equivalent of the specified number.
Syntax	<code>str(<i>approx_numeric</i> [, <i>length</i> [, <i>decimal</i>] ])</code>
Parameters	<p><i>approx_numeric</i> is any approximate numeric (float, real, or double precision) column name, variable, or constant expression.</p> <p><i>length</i> sets the number of characters to be returned (including the decimal point, all digits to the right and left of the decimal point, and blanks). The default is 10.</p> <p><i>decimal</i> sets the number of decimal digits to be returned. The default is 0.</p>

### Examples

#### Example 1

```
select str(1234.7, 4)
----
1235
```

#### Example 2

```
select str(-12345, 6)
-----
-12345
```

#### Example 3

```
select str(123.45, 5, 2)
-----
123.5
```

### Usage

- `str`, a string function, returns a character representation of the floating point number. For general information about string functions, see “String functions” on page 67.
- *length* and *decimal* are optional. If given, they must be nonnegative. `str` rounds the decimal portion of the number so that the results fit within the specified length. The length should be long enough to accommodate the decimal point and, if negative, the number’s sign. The decimal portion of the result is rounded to fit within the specified length. If the integer portion of the number does not fit within the length, however, `str` returns a row of asterisks of the specified length. For example:

```
select str(123.456, 2, 4)
```

--  
\*\*

A short *approx\_numeric* is right-justified in the specified length, and a long *approx\_numeric* is truncated to the specified number of decimal places.

- If *approx\_numeric* is NULL, returns NULL.

Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute str.
See also	<b>Functions</b> abs, ceiling, floor, round, sign

## str\_replace

Description	Replaces any instances of the second string expression ( <i>string_expression2</i> ) that occur within the first string expression ( <i>string_expression1</i> ) with a third expression ( <i>string_expression3</i> ).
Syntax	<code>replace("string_expression1", "string_expression2", "string_expression3")</code>
Parameters	<p><i>string_expression1</i> is the source string, or the string expression to be searched, expressed as char, varchar, unichar, univarchar, varbinary, or binary datatype.</p> <p><i>string_expression2</i> is the pattern string, or the string expression to find within the first expression (<i>string_expression1</i>). <i>string_expression2</i> is expressed as char, varchar, unichar, univarchar, varbinary, or binary datatype.</p> <p><i>string_expression3</i> is the replacement string expression, expressed as char, varchar, unichar, univarchar, binary, or varbinary datatype.</p>
Examples	<p><b>Example 1</b> Replaces the string <i>def</i> within the string <i>cdefghi</i> with <i>yyy</i>.</p> <pre>replace("cdefghi", "def", "yyy") ----- cyyyghi (1 row(s) affected)</pre> <p><b>Example 2</b> Replaces all spaces with "toyota".</p> <pre>select str_replace("chevy, ford, mercedes", " ", "toyota") ----- chevy,toyotaford,toyotamercedes (1 row(s) affected)</pre> <hr/> <p><b>Note</b> Adaptive Server converts an empty string constant to a string of one space automatically, to distinguish the string from NULL values.</p> <hr/> <p><b>Example 3</b> Returns "abcghijklm":</p> <pre>select str_replace("abcdefghijklm", "def", NULL) ----- abcghijklm (1 row affected)</pre>
Usage	<ul style="list-style-type: none"><li>Returns varchar data if <i>string_expression</i> (1, 2, or 3) is char or varchar.</li></ul>

- Returns univarchar data if *string\_expression* (1, 2, or 3) is unichar or univarchar.
- Returns varbinary data if *string\_expression* (1, 2, or 3) is binary or varbinary.
- All arguments must share the same datatype.
- If any of the three arguments is NULL, the function returns null.

str\_replace accepts NULL in the third parameter and treats it as an attempt to replace *string\_expression2* with NULL, effectively turning str\_replace into a “string cut” operation.

For example, the following returns “abcghijklm”:

```
str_replace("abcdefghijklm", "def", NULL)
```

- The result length may vary, depending upon what is known about the argument values when the expression is compiled. If all arguments are variables with known constant values, Adaptive Server calculates the result length as:

$$\text{result\_length} = ((s/p) * (r-p) + s)$$

where

s = length of source string

p = length of pattern string

r = length of replacement string

if (r-p) <= 0, result length = s

- If the source string (*string\_expression1*) is a column, and *string\_expression2* and *string\_expression3* are constant values known at compile time, Adaptive Server calculates the result length using the formula above.
- If Adaptive Server cannot calculate the result length because the argument values are unknown when the expression is compiled, the result length used is 255, unless traceflag 244 is on. In that case, the result length is 16384.
- result\_len never exceeds 16384.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute str\_replace.

See also

**Datatypes** char, varchar, binary, varbinary, unichar, univarchar

**Function** length

## stuff

Description	Returns the string formed by deleting a specified number of characters from one string and replacing them with another string.
Syntax	<code>stuff(char_expr1   uchar_expr1, start, length, char_expr2   uchar_expr2)</code>
Parameters	<p><i>char_expr1</i> is a character-type column name, variable, or constant expression of char, varchar, nchar, or nvarchar type.</p> <p><i>uchar_expr1</i> is a character-type column name, variable, or constant expression of unichar or univarchar type.</p> <p><i>start</i> specifies the character position at which to begin deleting characters.</p> <p><i>length</i> specifies the number of characters to delete.</p> <p><i>char_expr2</i> is another character-type column name, variable, or constant expression of char, varchar, nchar, or nvarchar type.</p> <p><i>uchar_expr2</i> is another character-type column name, variable, or constant expression of unichar or univarchar type.</p>
Examples	<p><b>Example 1</b></p> <pre>select stuff("abc", 2, 3, "xyz") ----- axyz</pre> <p><b>Example 2</b></p> <pre>select stuff("abcdef", 2, 3, null) go --- aef</pre> <p><b>Example 3</b></p> <pre>select stuff("abcdef", 2, 3, "") ----- a ef</pre>



---

Usage	<ul style="list-style-type: none"><li>• <code>stuff</code>, a string function, deletes <i>length</i> characters from <i>char_expr1</i> or <i>uchar_expr1</i> at <i>start</i>, then inserts <i>char_expr2</i> or <i>uchar_expr2</i> into <i>char_expr1</i> or <i>uchar_expr2</i> at <i>start</i>. For general information about string functions, see “String functions” on page 67.</li><li>• If the start position or the length is negative, a NULL string is returned. If the start position is longer than <i>expr1</i>, a NULL string is returned. If the length to be deleted is longer than <i>expr1</i>, <i>expr1</i> is deleted through its last character (see Example 1).</li><li>• If the start position falls in the middle of a surrogate pair, start is adjusted to be one less. If the start length position falls in the middle of a surrogate pair, length is adjusted to be one less.</li><li>• To use <code>stuff</code> to delete a character, replace <i>expr2</i> with NULL rather than with empty quotation marks. Using “ ” to specify a null character replaces it with a space (see Examples 2 and 3).</li><li>• If <i>char_expr1</i> or <i>uchar_expr1</i> is NULL, <code>stuff</code> returns NULL. If <i>char_expr1</i> or <i>uchar_expr1</i> is a string value and <i>char_expr2</i> or <i>uchar_expr2</i> is NULL, <code>stuff</code> replaces the deleted characters with nothing.</li><li>• If you give a varchar expression as one parameter and a unichar expression as the other, the varchar expression is implicitly converted to unichar (with possible truncation).</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>stuff</code> .
See also	<b>Functions</b> replicate, substring

## substring

Description	Returns the string formed by extracting the specified number of characters from another string.
Syntax	<code>substring(expression, start, length)</code>
Parameters	<p><i>expression</i> is a binary or character column name, variable, or constant expression. Can be char, nchar, unichar, varchar, univarchar, or nvarchar data, binary, or varbinary.</p> <p><i>start</i> specifies the character position at which the substring begins.</p> <p><i>length</i> specifies the number of characters in the substring.</p>
Examples	<p><b>Example 1</b> Displays the last name and first initial of each author, for example, “Bennet A.”:</p> <pre>select au_lname, substring(au_fname, 1, 1) from authors</pre> <p><b>Example 2</b> Converts the author’s last name to uppercase, then displays the first three characters:</p> <pre>select substring(upper(au_lname), 1, 3) from authors</pre> <p><b>Example 3</b> Concatenates pub_id and title_id, then displays the first six characters of the resulting string:</p> <pre>select substring((pub_id + title_id), 1, 6) from titles</pre> <p><b>Example 4</b> Extracts the lower four digits from a binary field, where each position represents two binary digits:</p> <pre>select substring(xactid, 5, 2) from syslogs</pre>
Usage	<ul style="list-style-type: none"><li>• substring, a string function, returns part of a character or binary string. For general information about string functions, see “String functions” on page 67.</li><li>• If substring’s second argument is NULL, the result is NULL. If substring’s first or third argument is NULL, the result is blank..</li></ul>

- If the start position from the beginning of *uchar\_expr1* falls in the middle of a surrogate pair, *start* is adjusted to one less. If the start length position from the beginning of *uchar\_expr1* falls in the middle of a surrogate pair, *length* is adjusted to one less.

Standards ANSI SQL – Compliance level: Transact-SQL extension.

Permissions Any user can execute substring.

See also **Functions** charindex, patindex, stuff

## sum

Description	Returns the total of the values.
Syntax	<code>sum([all   distinct] <i>expression</i>)</code>
Parameters	<p><b>all</b> applies sum to all values. <b>all</b> is the default.</p> <p><b>distinct</b> eliminates duplicate values before sum is applied. <b>distinct</b> is optional.</p> <p><b><i>expression</i></b> is a column name, constant, function, any combination of column names, constants, and functions connected by arithmetic or bitwise operators, or a subquery. With aggregates, an expression is usually a column name. For more information, see “Expressions” on page 275.</p>
Examples	<p><b>Example 1</b> Calculates the average advance and the sum of total sales for all business books. Each of these aggregate functions produces a single summary value for all of the retrieved rows:</p> <pre>select avg(advance), sum(total_sales) from titles where type = "business"</pre> <p><b>Example 2</b> Used with a group by clause, the aggregate functions produce single values for each group, rather than for the entire table. This statement produces summary values for each type of book:</p> <pre>select type, avg(advance), sum(total_sales) from titles group by type</pre> <p><b>Example 3</b> Groups the titles table by publishers, and includes only those groups of publishers who have paid more than \$25,000 in total advances and whose books average more than \$15 in price:</p> <pre>select pub_id, sum(advance), avg(price) from titles group by pub_id having sum(advance) &gt; \$25000 and avg(price) &gt; \$15</pre>
Usage	<ul style="list-style-type: none"><li>• <b>sum</b>, an aggregate function, finds the sum of all the values in a column. <b>sum</b> can only be used on numeric (integer, floating point, or money) datatypes. Null values are ignored in calculating sums.</li><li>• For general information about aggregate functions, see “Aggregate functions” on page 49.</li></ul>

- When you sum integer data, Adaptive Server treats the result as an int value, even if the datatype of the column is smallint or tinyint. When you sum bigint data, Adaptive Server treats the result as a bigint. To avoid overflow errors in DB-Library programs, declare all variables for results of averages or sums appropriately.
- You cannot use sum with the binary datatypes.
- This function defines only numeric types; use with Unicode expressions generates an error.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute sum.

See also

**Commands** compute clause, group by and having clauses, select, where clause

**Functions** count, max, min

## suser\_id

Description Returns the server user's ID number from the syslogins table.

Syntax `suser_id([server_user_name])`

Parameters *server\_user\_name*  
is an Adaptive Server login name.

Examples

**Example 1**

```
select suser_id()
-----
1
```

**Example 2**

```
select suser_id("margaret")
-----
5
```

Usage

- `suser_id`, a system function, returns the server user's ID number from `syslogins`. For general information about system functions, see "System functions" on page 68.
- To find the user's ID in a specific database from the `sysusers` table, use the `user_id` system function.
- If no *server\_user\_name* is supplied, `suser_id` returns the server ID of the current user.

Standards

ANSI SQL – Compliance level: Transact-SQL extension.

Permissions

Any user can execute `suser_id`.

See also

**Functions** `suser_name`, `user_id`

## suser\_name

Description	Returns the name of the current server user or the user whose server ID is specified.
Syntax	<code>suser_name([server_user_id])</code>
Parameters	<code>server_user_id</code> is an Adaptive Server user ID.
Examples	<p><b>Example 1</b></p> <pre>select suser_name() ----- sa</pre> <p><b>Example 2</b></p> <pre>select suser_name(4) ----- margaret</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>suser_name</code>, a system function, returns the server user's name. Server user IDs are stored in syslogins. If no <code>server_user_id</code> is supplied, <code>suser_name</code> returns the name of the current user.</li> <li>• For general information about system functions, see "System functions" on page 68.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>suser_name</code> .
See also	<b>Functions</b> <code>suser_id</code> , <code>user_name</code>

## syb\_quit

Description	Terminates the connection.
Syntax	syb_quit()
Examples	Terminates the connection in which the function is executed and returns an error message. <pre>select syb_quit() ----- CT-LIBRARY error:   ct_results(): network packet layer: internal net library error: Net-Library operation terminated due to disconnect</pre>
Usage	You can use syb_quit to terminate a script if the isql preprocessor command exit causes an error.
Permissions	Any user can execute syb_quit.



## syb\_sendmsg

Description	<b>UNIX only</b> Sends a message to a User Datagram Protocol (UDP) port.
Syntax	<code>syb_sendmsg ip_address, port_number, message</code>
Parameters	<p><i>ip_address</i> is the IP address of the machine where the UDP application is running.</p> <p><i>port_number</i> is the port number of the UDP port.</p> <p><i>message</i> is the message to send. It can be up to 255 characters in length.</p>
Examples	<p><b>Example 1</b> Sends the message “Hello” to port 3456 at IP address 120.10.20.5:</p> <pre>select syb_sendmsg("120.10.20.5", 3456, "Hello")</pre> <p><b>Example 2</b> Reads the IP address and port number from a user table, and uses a variable for the message to be sent:</p> <pre>declare @msg varchar(255) select @msg = "Message to send" select syb_sendmsg (ip_address, portnum, @msg) from sendports where username = user_name()</pre>
Usage	<ul style="list-style-type: none"> <li>• To enable the use of UDP messaging, a System Security Officer must set the configuration parameter <code>allow_sendmsg</code> to 1.</li> <li>• No security checks are performed with <code>syb_sendmsg</code>. Sybase strongly recommends that you do not use <code>syb_sendmsg</code> to send sensitive information across the network. By enabling this functionality, the user accepts any security problems that result from its use.</li> <li>• For a sample C program that creates a UDP port, see <code>sp_sendmsg</code>.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>syb_sendmsg</code> .
See also	<b>System procedure</b> <code>sp_sendmsg</code>

## tan

Description	Returns the tangent of the specified angle (in radians).
Syntax	<code>tan(<i>angle</i>)</code>
Parameters	<i>angle</i> is the size of the angle in radians, expressed as a column name, variable, or expression of type float, real, double precision, or any datatype that can be implicitly converted to one of these types.
Examples	<pre>select tan(60) -----           0.320040</pre>
Usage	<ul style="list-style-type: none"><li>tan, a mathematical function, returns the tangent of the specified angle (measured in radians).</li><li>For general information about mathematical functions, see “Mathematical functions” on page 65.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute tan.
See also	<b>Functions</b> atan, atn2, degrees, radians

## tempdb\_id

Description	Reports the temporary database to which a given session is assigned. The input of the tempdb_id function is a server process ID, and its output is the temporary database to which the process is assigned. If you do not provide a server process, tempdb_id reports the dbid of the temporary database assigned to the current process.
Syntax	tempdb_id()
Examples	Finds all the server processes that are assigned to a given temporary database: <pre>select spid from master..sysprocesses       where tempdb_id(spid) = db_id("tempdatabase")</pre>
Usage	select tempdb_id() gives the same result as select @@tempdbid.
See also	<b>Commands</b> select

## textptr

**Description** Returns a pointer to the first page of a text, image, or unitext column.

**Syntax** textptr(*column\_name*)

**Parameters** *column\_name*  
is the name of a text column.

**Examples** **Example 1** Uses the textptr function to locate the text column, copy, associated with au\_id 486-29-1786 in the author's blurbs table. The text pointer is placed in local variable @val and supplied as a parameter to the readtext command, which returns 5 bytes, starting at the second byte (offset of 1):

```
declare @val binary(16)
select @val = textptr(copy) from blurbs
where au_id = "486-29-1786"
readtext blurbs.copy @val 1 5
```

**Example 2** Selects the title\_id column and the 16-byte text pointer of the copy column from the blurbs table:

```
select au_id, textptr(copy) from blurbs
```

**Usage**

- textptr, a text and image function, returns the text pointer value, a 16-byte varbinary value.
- If a text, unitext, or image column has not been initialized by a non-null insert or by any update statement, textptr returns a NULL pointer. Use textvalid to check whether a text pointer exists. You cannot use writetext or readtext without a valid text pointer.
- For general information about text and image functions, see “Text and image functions” on page 69.

---

**Note** Trailing  $\text{f}$  in varbinary values are truncated when the values are stored in tables. If you are storing text pointer values in a table, use binary as the datatype for the column.

---

**Standards** ANSI SQL – Compliance level: Transact-SQL extension.

**Permissions** Any user can execute textptr.

**See also** **Datatypes** text, image, and unitext datatypes

**Function** textvalid

**Commands** insert, update, readtext, writetext

## textvalid

Description	Returns 1 if the pointer to the specified text or unitext column is valid; 0 if it is not.
Syntax	<code>textvalid("table_name.column_name", textpointer)</code>
Parameters	<p><i>table_name.column_name</i> is the name of a table and its text column.</p> <p><i>textpointer</i> is a text pointer value.</p>
Examples	<p>Reports whether a valid text pointer exists for each value in the blurb column of the texttest table:</p> <pre>select textvalid ("textttest.blurb", textptr(blurb)) from textttest</pre>
Usage	<ul style="list-style-type: none"> <li>• textvalid, a text and image function, checks that a given text pointer is valid. Returns 1 if the pointer is valid, or 0 if it is not.</li> <li>• The identifier for a text or an image column must include the table name.</li> <li>• For general information about text and image functions, see “Text and image functions” on page 69.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute textvalid.
See also	<p><b>Datatypes</b> text, image, and unitext datatypes</p> <p><b>Function</b> textptr</p>

## to\_unichar

Description	Returns a unichar expression having the value of the integer expression.
Syntax	to_unichar ( <i>integer_expr</i> )
Parameters	<i>integer_expr</i> is any integer (tinyint, smallint, or int) column name, variable, or constant expression.
Usage	<ul style="list-style-type: none"><li>• to_unichar, a string function, converts a Unicode integer value to a Unicode character value.</li><li>• If a unichar expression refers to only half of a surrogate pair, an error message appears and the operation is aborted.</li><li>• If a <i>integer_expr</i> is NULL, to_unichar returns NULL.</li><li>• For general information about string functions, see “String functions” on page 67.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute to_unichar.
See also	<b>Datatypes</b> text, image, and unitext datatypes <b>Function</b> char

## tran\_dumpable\_status

**Description** Returns a true/false indication of whether dump transaction is allowed.

**Syntax** `tran_dumpable_status("database_name")`

**Parameters** *database\_name*  
is the name of the target database.

**Examples** Checks to see if the pubs2 database can be dumped:

```
1> select tran_dumpable_status("pubs2")
2> go

-----
                106

(1 row affected)
```

In this example, you cannot dump pubs2. The return code of 106 is a sum of all the conditions met (2, 8, 32, 64). See the Usage section for a description of the return codes.

**Usage** `tran_dumpable_status` allows you to determine if dump transaction is allowed on a database without having to run the command. `tran_dumpable_status` performs all of the checks that Adaptive Server performs when dump transaction is issued.

If `tran_dumpable_status` returns 0, you can perform the dump transaction command on the database. If it returns any other value, it cannot. The non-0 values are:

- 1 – A database with the name you specified does not exist.
- 2 – A log does not exist on a separate device.
- 4 – The log first page is in the bounds of a data-only disk fragment.
- 8 – the trunc log on chkpt option is set for the database.
- 16 – Non-logged writes have occurred on the database.
- 32 – Truncate-only dump tran has interrupted any coherent sequence of dumps to dump devices.
- 64 – Database is newly created or upgraded. Transaction log may not be dumped until a dump database has been performed.

**Standards** ANSI SQL – Compliance level: Transact-SQL extension.

**Permissions** Any user can execute this function.

**See also** **Command** dump transaction

## tsequal

**Description** Compares timestamp values to prevent update on a row that has been modified since it was selected for browsing.

**Syntax** `tsequal(browsed_row_timestamp, stored_row_timestamp)`

**Parameters**

*browsed\_row\_timestamp*  
is the timestamp column of the browsed row.

*stored\_row\_timestamp*  
is the timestamp column of the stored row.

**Examples** Retrieves the timestamp column from the current version of the publishers table and compares it to the value in the timestamp column that has been saved. If the values in the two timestamp columns are equal, tsequal updates the row. If the values are not equal, tsequal returns this error message:

```
update publishers
set city = "Springfield"
where pub_id = "0736"
and tsequal(timestamp, 0x0001000000002ea8)
```

**Usage**

- tsequal, a system function, compares the timestamp column values to prevent an update on a row that has been modified since it was selected for browsing. For general information about system functions, see “System functions” on page 68.
- tsequal allows you to use browse mode without calling the dbqual function in DB-Library. Browse mode supports the ability to perform updates while viewing data. It is used in front-end applications using Open Client and a host programming language. A table can be browsed if its rows have been timestamped.
- To browse a table in a front-end application, append the for browse keywords to the end of the select statement sent to Adaptive Server. For example:

```
Start of select statement in an Open Client application
...
for browse
```

*Completion of the Open Client application routine*

- Do not use tsequal in the where clause of a select statement; only in the where clause of insert and update statements where the rest of the where clause matches a single unique row.



If you use a timestamp column as a search clause, compare it like a regular varbinary column; that is, `timestamp1 = timestamp2`.

#### Timestamping a new table for browsing

- When creating a new table for browsing, include a column named `timestamp` in the table definition. The column is automatically assigned a datatype of `timestamp`; you do not have to specify its datatype. For example:

```
create table newtable(col1 int, timestamp, col3 char(7))
```

Whenever you insert or update a row, Adaptive Server timestamps it by automatically assigning a unique varbinary value to the `timestamp` column.

#### Timestamping an existing table

- To prepare an existing table for browsing, add a column named `timestamp` using `alter table`. For example, to add a `timestamp` column with a `NULL` value to each existing row:

```
alter table oldtable add timestamp
```

To generate a timestamp, update each existing row without specifying new column values:

```
update oldtable
set col1 = col1
```

Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>tsequal</code> .
See also	<b>Datatype</b> Timestamp datatype

## uhighsurr

Description	Returns 1 if the Unicode value at position <i>start</i> is the high half of a surrogate pair (which should appear first in the pair). Returns 0 otherwise.
Syntax	uhighsurr( <i>uchar_expr</i> , <i>start</i> )
Parameters	<i>uchar_expr</i> is a character-type column name, variable, or constant expression of unichar or univarchar type.  <i>start</i> specifies the character position to investigate.
Usage	<ul style="list-style-type: none"><li>• uhighsurr, a string function, allows you to write explicit code for surrogate handling. Specifically, if a substring starts on a Unicode character where uhighsurr is true, extract a substring of at least 2 Unicode values (<i>substr</i> does not extract half of a surrogate pair).</li><li>• If <i>uchar_expr</i> is NULL, uhighsurr returns NULL.</li><li>• For general information about string functions, see “String functions” on page 67.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute uhighsurr.
See also	<b>Function</b> ulowsurr

## ulowsurr

Description	Returns 1 if the Unicode value at position <i>start</i> is the low half of a surrogate pair (which should appear second in the pair). Returns 0 otherwise.
Syntax	<code>ulowsurr(<i>uchar_expr</i>, start)</code>
Parameters	<i>uchar_expr</i> is a character-type column name, variable, or constant expression of unichar or univarchar type.  <i>start</i> specifies the character position to investigate.
Usage	<ul style="list-style-type: none"><li>• <code>ulowsurr</code>, a string function, allows you to write explicit code around adjustments performed by <code>substr()</code>, <code>stuff()</code>, and <code>right()</code>. Specifically, if a substring ends on a Unicode value where <code>ulowsurr()</code> is true, the user knows to extract a substring of 1 less characters (or 1 more). <code>substr()</code> does not extract a string that contains an unmatched surrogate pair.</li><li>• If <i>uchar_expr</i> is NULL, <code>ulowsurr</code> returns NULL.</li><li>• For general information about string functions, see “String functions” on page 67.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>ulowsurr</code> .
See also	<b>Function</b> <code>uhighsurr</code>

## upper

Description	Returns the uppercase equivalent of the specified string.
Syntax	<code>upper(char_expr)</code>
Parameters	<i>char_expr</i> is a character-type column name, variable, or constant expression of char, unichar, varchar, nchar, nvarchar, or univarchar type.
Examples	<pre>select upper("abcd") ----- ABCD</pre>
Usage	<ul style="list-style-type: none"><li>• upper, a string function, converts lowercase to uppercase, returning a character value.</li><li>• If <i>char_expr</i> or <i>uchar_expr</i> is NULL, upper returns NULL.</li><li>• Characters that have no upper-case equivalent are left unmodified.</li><li>• If a unichar expression is created containing only half of a surrogate pair, an error message appears and the operation is aborted.</li><li>• For general information about string functions, see “String functions” on page 67.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute upper.
See also	<b>Function</b> lower

## uscalar

Description	Returns the Unicode scalar value for the first Unicode character in an expression.
Syntax	<code>uscalar(<i>uchar_expr</i>)</code>
Parameters	<i>uchar_expr</i> is a character-type column name, variable, or constant expression of <code>unichar</code> , or <code>univarchar</code> type.
Usage	<ul style="list-style-type: none"><li>• <code>uscalar</code>, a string function, returns the Unicode value for the first Unicode character in an expression.</li><li>• If <i>uchar_expr</i> is <code>NULL</code>, returns <code>NULL</code>.</li><li>• If <code>uscalar</code> is called on a <i>uchar_expr</i> containing an unmatched surrogate half, and error occurs and the operation is aborted.</li><li>• For general information about string functions, see “String functions” on page 67.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>uscalar</code> .
See also	<b>Functions</b> <code>ascii</code>

## used\_pages

Description	Reports the number of pages used by a table, an index, or a specific partition. Unlike <code>data_pages</code> , <code>used_pages</code> does include pages used for internal structures. This function replaces the <code>used_pgs</code> function used in versions of Adaptive Server earlier than 15.0.
Syntax	<code>used_pages(<i>dbid</i>, <i>object_id</i> [, <i>indid</i> [, <i>ptnid</i>]])</code>
Parameters	<p><i>dbid</i> the database id where target object resides.</p> <p><i>object_id</i> is the object ID of the table for which you want to see the used pages. To see the pages used by an index, specify the object ID of the table to which the index belongs.</p> <p><i>indid</i> is the index id of interest.</p> <p><i>ptnid</i> is the partition id of interest.</p>
Examples	<p><b>Example 1</b> Returns the number of pages used by the object with a object ID of 31000114 in the specified database (including any indexes):</p> <pre>select used_pages(5, 31000114)</pre> <p><b>Example 2</b> Returns the number of pages used by the object in the data layer, regardless of whether or not a clustered index exists:</p> <pre>select used_pages(5, 31000114, 0)</pre> <p><b>Example 3</b> Returns the number of pages used by the object in the index layer for a clustered index. This does not include the pages used by the data layer:</p> <pre>select used_pages(5, 31000114, 1)</pre> <p><b>Example 4</b> Returns the number of pages used by the object in the data layer of the specific partition, which in this case is 2323242432:</p> <pre>select used_pages(5, 31000114, 0, 2323242432)</pre>
Usage	<ul style="list-style-type: none"><li>• <code>used_pages(<i>dbid</i>, <i>objid</i>, 0)</code> is identical to <code>used_pages(<i>dbid</i>, <i>objid</i>, 1)</code> in the case of an all-page lock table with a clustered index. This is similar to the old <code>used_pgs(<i>objid</i>, <i>doampg</i>, <i>ioampg</i>)</code> function.</li><li>• All erroneous conditions result in a return value of zero.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>used_pgs</code> .

See also **Functions** `data_pages`, `object_id`

## user

Description Returns the name of the current user.

Syntax user

Parameters None.

Examples

```
select user
-----
dbo
```

- Usage
- user, a system function, returns the user’s name.
  - If the sa\_role is active, you are automatically the Database Owner in any database you are using. Inside a database, the user name of the Database Owner is always “dbo”.
  - For general information about system functions, see “System functions” on page 68.

Standards ANSI SQL – Compliance level: Transact-SQL extension.

Permissions Any user can execute user.

See also **Functions** user\_name



## user\_id

Description	Returns the ID number of the specified user or of the current user in the database.
Syntax	<code>user_id([user_name])</code>
Parameters	<i>user_name</i> is the name of the user.
Examples	<p><b>Example 1</b></p> <pre>select user_id() ----- 1</pre> <p><b>Example 2</b></p> <pre>select user_id("margaret") ----- 4</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>user_id</code>, a system function, returns the user's ID number. For general information about system functions, see "System functions" on page 68.</li> <li>• <code>user_id</code> reports the number from <code>sysusers</code> in the current database. If no <i>user_name</i> is supplied, <code>user_id</code> returns the ID of the current user. To find the server user ID, which is the same number in every database on Adaptive Server, use <code>suser_id</code>.</li> <li>• Inside a database, the "guest" user ID is always 2.</li> <li>• Inside a database, the <code>user_id</code> of the Database Owner is always 1. If you have the <code>sa_role</code> active, you are automatically the Database Owner in any database you are using. To return to your actual user ID, use <code>set sa_role off</code> before executing <code>user_id</code>. If you are not a valid user in the database, Adaptive Server returns an error when you use <code>set sa_role off</code>.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	You must System Administrator or System Security Officer to use this function on a <i>user_name</i> other than your own.
See also	<p><b>Commands</b> <code>setuser</code></p> <p><b>Functions</b> <code>suser_id</code>, <code>user_name</code></p>

## **user\_name**

**Description** Returns the name within the database of the specified user or of the current user.

**Syntax** `user_name([user_id])`

**Parameters** *user\_id*  
is the ID of a user.

**Examples** **Example 1**

```
select user_name()  
-----  
dbo
```

**Example 2**

```
select user_name(4)  
-----  
margaret
```

**Usage**

- `user_name`, a system function, returns the user’s name, based on the user’s ID in the current database. For general information about system functions, see “System functions” on page 68.
- If no *user\_id* is supplied, `user_name` returns the name of the current user.
- If the `sa_role` is active, you are automatically the Database Owner in any database you are using. Inside a database, the `user_name` of the Database Owner is always “dbo”.

**Standards** ANSI SQL – Compliance level: Transact-SQL extension.

**Permissions** You must be a System Administrator or System Security Officer to use this function on a *user\_id* other than your own.

**See also** **Functions** `suser_name`, `user_id`

## valid\_name

Description	Returns 0 if the specified string is not a valid identifier or a number other than 0 if the string is a valid identifier, and can be up to 255 bytes in length.
Syntax	<code>valid_name(character_expression [, maximum_length])</code>
Parameters	<p><i>character_expression</i> is a character-type column name, variable, or constant expression of char, varchar, nchar or nvarchar type. Constant expressions must be enclosed in quotation marks.</p> <p><i>maximum_length</i> is an integer larger than 0 and less than or equal to 255. The default value is 30. If the identifier length is larger than the second argument, <code>valid_name</code> returns 0, and returns a value greater than zero if the identifier length is invalid.</p>
Examples	<p>Creates a procedure to verify that identifiers are valid:</p> <pre>create procedure chkname @name varchar(30) as     if valid_name(@name) = 0     print "name not valid"</pre>
Usage	<ul style="list-style-type: none"> <li>• <code>valid_name</code>, a system function, returns 0 if the <i>character_expression</i> is not a valid identifier (illegal characters, more than 30 bytes long, or a reserved word), or a number other than 0 if it is a valid identifier.</li> <li>• Adaptive Server identifiers can be a maximum of 16384 bytes in length, whether single-byte or multibyte characters are used. The first character of an identifier must be either an alphabetic character, as defined in the current character set, or the underscore (<code>_</code>) character. Temporary table names, which begin with the pound sign (<code>#</code>), and local variable names, which begin with the at sign (<code>@</code>), are exceptions to this rule. <code>valid_name</code> returns 0 for identifiers that begin with the pound sign (<code>#</code>) and the at sign (<code>@</code>).</li> <li>• For general information about system functions, see “System functions” on page 68.</li> </ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute <code>valid_name</code> .
See also	<b>System procedure</b> <code>sp_checkreswords</code>

## **valid\_user**

Description	Returns 1 if the specified ID is a valid user or alias in at least one database on this Adaptive Server.
Syntax	<code>valid_user(server_user_id)</code>
Parameters	<i>server_user_id</i> is a server user ID. Server user IDs are stored in the <code>suid</code> column of <code>syslogins</code> .
Examples	<pre>select valid_user(4) ----- 1</pre>
Usage	<ul style="list-style-type: none"><li>• <code>valid_user</code>, a system function, returns 1 if the specified ID is a valid user or alias in at least one database on this Adaptive Server.</li><li>• For general information about system functions, see “System functions” on page 68.</li></ul>
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	You must be a System Administrator or a System Security Officer to use this function on a <code>server_user_id</code> other than your own.
See also	<b>System procedures</b> <code>sp_addlogin</code> , <code>sp_adduser</code>

## year

Description	Returns an integer that represents the year in the datepart of a specified date.
Syntax	<code>year(date_expression)</code>
Parameters	<i>date_expression</i> is an expression of type datetime, smalldatetime, date, time or a character string in a datetime format.
Examples	Returns the integer 03: <pre>year ("11/02/03") ----- 03 (1 row(s) affected)</pre>
Usage	<code>year(date_expression)</code> is equivalent to <code>datepart(yy, date_expression)</code> .
Standards	ANSI SQL – Compliance level: Transact-SQL extension.
Permissions	Any user can execute year.
See also	<b>Datatypes</b> datetime, smalldatetime, date <b>Functions</b> datepart, day, month



Topics	Page
Adaptive Server global variables	267

## Adaptive Server global variables

Global variables are system-defined variables updated by Adaptive Server while the system is running. Some global variables are session-specific, while others are server instance-specific. For example, @@error contains the last error number generated by the system for a given user connection.

See `get_appcontext` and `set_appcontext` to specify application context variables.

To view the value for any global variable, enter:

```
select variable_name
```

For example:

```
select @@char_convert
```

Table 3-1 lists the global variables available for Adaptive Server:

**Table 3-1: Adaptive Server global variables**

Global variable	Definition
@@authmech	A read-only variable that indicates the mechanism used to authenticate the user.
@@bootcount	Returns the number of times an Adaptive Server installation has been booted.
@@boottime	Returns the date and time Adaptive Server was last booted.
@@bulkarraysize	Returns the number of rows to be buffered in local server memory before being transferred using the bulk copy interface. Used only with Component Integration Services for transferring rows to a remote server using <code>select into</code> . For more information, see the <i>Component Integration Services User's Guide</i> .
@@bulkbatchsize	Returns the number of rows transferred to a remote server via <code>select into proxy_table</code> using the bulk interface. Used only with Component Integration Services for transferring rows to a remote server using <code>select into</code> . For more information, see the <i>Component Integration Services User's Guide</i> .

Global variable	Definition
<code>@@char_convert</code>	Returns 0 if character set conversion is not in effect. Returns 1 if character set conversion is in effect.
<code>@@cis_rpc_handling</code>	Returns 0 if cis rpc handling is off. Returns 1 if cis rpc handling is on. For more information, see the <i>Component Integration Services User's Guide</i> .
<code>@@cis_version</code>	Returns the date and version of Component Integration Services.
<code>@@client_csexpansion</code>	Returns the expansion factor used when converting from the server character set to the client character set. For example, if it contains a value of 2, a character in the server character set could take up to twice the number of bytes after translation to the client character set.
<code>@@client_csid</code>	Returns -1 if the client character set has never been initialized. Returns the client character set ID from syscharsets for the connection if the client character set has been initialized.
<code>@@client_csname</code>	Returns NULL if client character set has never been initialized; returns the name of the character set for the connection if the client character set has been initialized.
<code>@@cmpstate</code>	Returns the current mode of Adaptive Server in a high availability environment.
<code>@@connections</code>	Returns the number of user logins attempted.
<code>@@cpu_busy</code>	Returns the number of seconds, in CPU time, that Adaptive Server's CPU was performing Adaptive Server work.
<code>@@cursor_rows</code>	<p>A global variable designed specifically for scrollable cursors. Displays the total number of rows in the cursor result set. Returns the following values: -1:</p> <ul style="list-style-type: none"> <li>The cursor is dynamic. Because dynamic cursors reflect all changes, the number of rows that qualify for the cursor is constantly changing. You can never be certain that all the qualified rows are retrieved.</li> <li>The cursor is semi_sensitive and scrollable, but the scrolling worktable is not yet fully populated. The number of rows that qualify the cursor is unknown at the time this value is retrieved.</li> </ul> <p>0: Either no cursors are open, no rows qualify for the last opened cursor, or the last open cursor is closed or deallocated.</p> <p>n: The last opened or fetched cursor result set is fully populated. The value returned is the total number of rows in the cursor result set.</p>
<code>@@curloid</code>	Either no cursors are open, no rows qualify for the last opened cursor, or the last open cursor is closed or deallocated.
<code>@@datefirst</code>	<p>Set using <code>set datefirst n</code> where <i>n</i> is a value between 1 and 7. Returns the current value of <code>@@datefirst</code>, indicating the specified first day of each week, expressed as tinyint.</p> <p>The default value in Adaptive Server is Sunday (based on the <code>us_language</code> default), which you set by specifying <code>set datefirst 7</code>. See the <code>datefirst</code> option of the <code>set</code> command for more information on settings and values.</p>
<code>@@dbts</code>	Returns the timestamp of the current database.
<code>@@error</code>	Returns the error number most recently generated by the system.



Global variable	Definition
<code>@@errorlog</code>	Returns the full path to the directory in which the Adaptive Server errorlog is kept, relative to <code>\$\$SYBASE</code> directory ( <code>%SYBASE%</code> on NT).
<code>@@failedoverconn</code>	Returns a value greater than 0 if the connection to the primary companion has failed over and is executing on the secondary companion server. Used only in a high availability environment, and is session-specific.
<code>@@fetch_status</code>	Returns values: 0: fetch operation successful; -1: fetch operation unsuccessful; -2: value reserved for future use.
<code>@@guestuserid</code>	Returns the ID of the guest user.
<code>@@hacmpservername</code>	Returns the name of the companion server in a high availability setup.
<code>@@haconnection</code>	Returns a value greater than 0 if the connection has the failover property enabled. This is a session-specific property.
<code>@@heapmemsize</code>	Returns the size of the heap memory pool, in bytes. See the <i>System Administration Guide</i> for more information on heap memory.
<code>@@identity</code>	Returns the most recently generated IDENTITY column value.
<code>@@idle</code>	Returns the number of seconds, in CPU time, that Adaptive Server has been idle.
<code>@@invaliduserid</code>	Returns a value of -1 for an invalid user ID.
<code>@@io_busy</code>	Returns the number of seconds in CPU time that Adaptive Server has spent doing input and output operations.
<code>@@isolation</code>	Returns the value of the session-specific isolation level (0, 1, or 3) of the current Transact-SQL program.
<code>@@kernel_addr</code>	Returns the starting address of the first shared memory region that contains the kernel region. The result is in the form of <i>Oxaddress pointer value</i> .
<code>@@kernel_size</code>	Returns the size of the kernel region that is part of the first shared memory region.
<code>@@langid</code>	Returns the server-wide language ID of the language in use, as specified in <code>syslanguages.langid</code> .
<code>@@language</code>	Returns the name of the language in use, as specified in <code>syslanguages.name</code> .
<code>@@lock_timeout</code>	Set using <code>set lock wait n</code> . Returns the current <code>lock_timeout</code> setting, in milliseconds. <code>@@lock_timeout</code> returns the value of <code>n</code> . The default value is no timeout. If no <code>set lock wait n</code> is executed at the beginning of the session, <code>@@lock_timeout</code> returns -1.
<code>@@maxcharlen</code>	Returns the maximum length, in bytes, of a character in Adaptive Server's default character set.
<code>@@max_connections</code>	Returns the maximum number of simultaneous connections that can be made with Adaptive Server in the current computer environment. You can configure Adaptive Server for any number of connections less than or equal to the value of <code>@@max_connections</code> with the number of user connections configuration parameter.
<code>@@maxgroupid</code>	Returns the highest group user ID. The highest value is 1048576.
<code>@@maxpagesize</code>	Returns the server's logical page size.
<code>@@max_precision</code>	Returns the precision level used by decimal and numeric datatypes set by the server. This value is a fixed constant of 38.
<code>@@maxspid</code>	Returns maximum valid value for the <code>spid</code> .

<b>Global variable</b>	<b>Definition</b>
<code>@@maxsuid</code>	Returns the highest server user ID. The default value is 2147483647.
<code>@@maxuserid</code>	Returns the highest user ID. The highest value is 2147483647.
<code>@@mempool_addr</code>	Returns the global memory pool table address. The result is in the form <i>0xaddress pointer value</i> . This variable is for internal use.
<code>@@min_poolsize</code>	Returns the minimum size of a named cache pool, in kilobytes. It is calculated based on the <code>DEFAULT_POOL_SIZE</code> , which is 256, and the current value of max database page size.
<code>@@mingroupid</code>	Returns the lowest group user ID. The lowest value is 16384.
<code>@@minspid</code>	Returns 1, which is the lowest value for spid.
<code>@@minsuid</code>	Returns the minimum server user ID. The lowest value is -32768.
<code>@@minuserid</code>	Returns the lowest user ID. The lowest value is -32768.
<code>@@monitors_active</code>	Reduces the number of messages displayed by <code>sp_sysmon</code> .
<code>@@ncharsize</code>	Returns the maximum length, in bytes, of a character set in the current server default character set.
<code>@@nestlevel</code>	Returns the current nesting level.
<code>@@nodeid</code>	Returns the current installation's 48-bit node identifier. Adaptive Server generates a nodeid the first time the master device is first used, and uniquely identifies an Adaptive Server installation.
<code>@@optgoal</code>	Returns the current optimization goal setting for query optimization
<code>@@options</code>	Returns a hexadecimal representation of the session's set options.
<code>@@opttimeout</code>	Returns the current optimization timeout limit setting for query optimization
<code>@@pack_received</code>	Returns the number of input packets read by Adaptive Server.
<code>@@pack_sent</code>	Returns the number of output packets written by Adaptive Server.
<code>@@packet_errors</code>	Returns the number of errors detected by Adaptive Server while reading and writing packets.
<code>@@pagesize</code>	Returns the server's virtual page size.
<code>@@parallel_degree</code>	Returns the current maximum parallel degree setting.
<code>@@probesuid</code>	Returns a value of 2 for the probe user ID.
<code>@@procid</code>	Returns the stored procedure ID of the currently executing procedure.

Global variable	Definition
<code>@@recovery_state</code>	<p>Indicates whether Adaptive Server is in recovery based on these returns:</p> <ul style="list-style-type: none"> <li>• <code>NOT_IN_RECOVERY</code> – Adaptive Server is not in startup recovery or in failover recovery. Recovery has been completed and all databases that can be online are brought online.</li> <li>• <code>RECOVERY_TUNING</code> – Adaptive Server is in recovery (either startup or failover) and is tuning the optimal number of recovery tasks.</li> <li>• <code>BOOTIME_RECOVERY</code> – Adaptive Server is in startup recovery and has completed tuning the optimal number of tasks. Not all databases have been recovered.</li> <li>• <code>FAILOVER_RECOVER</code> – Adaptive Server is in recovery during an HA failover and has completed tuning the optimal number of recovery tasks. All databases are not brought online yet.</li> </ul>
<code>@@repartition_degree</code>	Returns the current dynamic repartitioning degree setting
<code>@@resource_granularity</code>	Returns the maximum resource usage hint setting for query optimization
<code>@@rowcount</code>	<p>The value of <code>@@rowcount</code> is affected by whether the specified cursor is forward-only or scrollable.</p> <p>If the cursor is the default, non-scrollable cursor, the value of <code>@@rowcount</code> increments one by one, in the forward direction only, until the number of rows in the result set are fetched. These rows are fetched from the underlying tables to the client. The maximum value for <code>@@rowcount</code> is the number of rows in the result set.</p> <p><i>In the default cursor, <code>@@rowcount</code> is set to 0 by any command that does not return or affect rows, such as an if or set command, or an update or delete statement that does not affect any rows.</i></p> <p>If the cursor is scrollable, there is no maximum value for <code>@@rowcount</code>. The value continues to increment with each fetch, regardless of direction, and there is no maximum value. The <code>@@rowcount</code> value in scrollable cursors reflects the number of rows fetched from the result set, not from the underlying tables, to the client.</p>
<code>@@scan_parallel_degree</code>	Returns the current maximum parallel degree setting for nonclustered index scans.
<code>@@servername</code>	Returns the name of Adaptive Server.
<code>@@setrowcount</code>	Returns the current value for set rowcount
<code>@@shmem_flags</code>	Returns the shared memory region properties. This variable is for internal use. There are a total of 13 different properties values corresponding to 13 bits in the integer. The valid values represented from low to high bit are: <code>MR_SHARED</code> , <code>MR_SPECIAL</code> , <code>MR_PRIVATE</code> , <code>MR_READABLE</code> , <code>MR_WRITABLE</code> , <code>MR_EXECUTABLE</code> , <code>MR_HWCOHERENCY</code> , <code>MR_SWCOHERENC</code> , <code>MR_EXACT</code> , <code>MR_BEST</code> , <code>MR_NAIL</code> , <code>MR_PSUEDO</code> , <code>MR_ZERO</code> .
<code>@@spid</code>	Returns the server process ID of the current process.
<code>@@sqlstatus</code>	Returns status information (warning exceptions) resulting from the execution of a fetch statement.

Global variable	Definition
<code>@@ssl_ciphersuite</code>	Returns NULL if SSL is not used on the current connection; otherwise, it returns the name of the cipher suite you chose during the SSL handshake on the current connection.
<code>@@stringsize</code>	Returns the amount of character data returned from a <code>toString()</code> method. The default is 50. Max values may be up to 2GB. A value of zero specifies the default value. See the <i>Component Integration Services User's Guide</i> for more information.
<code>@@tempdbid</code>	Returns a valid temporary database ID (dbid) of the session's assigned temporary database.
<code>@@textcolid</code>	Returns the column ID of the column referenced by <code>@@textptr</code> .
<code>@@textdataptnid</code>	Returns the partition ID of a text partition containing the column referenced by <code>@@textptr</code> .
<code>@@textdbid</code>	Returns the database ID of a database containing an object with the column referenced by <code>@@textptr</code> .
<code>@@textobjid</code>	Returns the object ID of an object containing the column referenced by <code>@@textptr</code> .
<code>@@textptnid</code>	Returns the partition ID of a data partition containing the column referenced by <code>@@textptr</code> .
<code>@@textptr</code>	Returns the text pointer of the last text, unitext, or image column inserted or updated by a process (Not the same as the <code>textptr</code> function).
<code>@@textptr_parameters</code>	Returns 0 if the current status of the <code>textptr_parameters</code> configuration parameter is off. Returns 1 if the current status of the <code>textptr_parameters</code> is on. See the <i>Component Integration Services User's Guide</i> for more information.
<code>@@textsize</code>	Returns the limit on the number of bytes of text, unitext, or image data a <code>select</code> returns. Default limit is 32K bytes for <code>isql</code> ; the default depends on the client software. Can be changed for a session with <code>set textsize</code> .
<code>@@textts</code>	Returns the text timestamp of the column referenced by <code>@@textptr</code> .
<code>@@thresh_hysteresis</code>	Returns the decrease in free space required to activate a threshold. This amount, also known as the hysteresis value, is measured in 2K database pages. It determines how closely thresholds can be placed on a database segment.
<code>@@timeticks</code>	Returns the number of microseconds per tick. The amount of time per tick is machine-dependent.
<code>@@total_errors</code>	Returns the number of errors detected by Adaptive Server while reading and writing.
<code>@@total_read</code>	Returns the number of disk reads by Adaptive Server.
<code>@@total_write</code>	Returns the number of disk writes by Adaptive Server.
<code>@@tranchained</code>	Returns 0 if the current transaction mode of the Transact-SQL program is unchained. Returns 1 if the current transaction mode of the Transact-SQL program is chained.
<code>@@trancount</code>	Returns the nesting level of transactions in the current user session.
<code>@@transactional_rpc</code>	Returns 0 if RPCs to remote servers are transactional. Returns 1 if RPCs to remote servers are not transactional. For more information, see <code>enable xact coordination</code> and <code>set option transactional_rpc</code> in the <i>Reference Manual</i> . Also, see the <i>Component Integration Services User's Guide</i> .

---

<b>Global variable</b>	<b>Definition</b>
<i>@@transtate</i>	Returns the current state of a transaction after a statement executes in the current user session.
<i>@@unicharsize</i>	Returns 2, the size of a character in unichar.
<i>@@version</i>	Returns the date, version string, and so on of the current release of Adaptive Server.
<i>@@version_number</i>	Returns the whole version of the current release of Adaptive Server as an integer
<i>@@version_as_integer</i>	Returns the number of the last upgrade version of the current release of Adaptive Server as an integer. For example, <i>@@version_as_integer</i> returns 12500 if you are running Adaptive Server version 12.5, 12.5.0.3, or 12.5.1.



# Expressions, Identifiers, and Wildcard Characters

This chapter describes Transact-SQL expressions, valid identifiers, and wildcard characters.

Topics covered are:

Topics	Page
Expressions	275
Identifiers	285
Pattern matching with wildcard characters	293

## Expressions

An expression is a combination of one or more constants, literals, functions, column identifiers and/or variables, separated by operators, that returns a single value. Expressions can be of several types, including **arithmetic**, **relational**, **logical** (or **Boolean**), and **character string**. In some Transact-SQL clauses, a subquery can be used in an expression. A case expression can be used in an expression.

Table 4-1 lists the types of expressions that are used in Adaptive Server syntax statements.

**Table 4-1: Types of expressions used in syntax statements**

Usage	Definition
expression	Can include constants, literals, functions, column identifiers, variables, or parameters
logical expression	An expression that returns TRUE, FALSE, or UNKNOWN
constant expression	An expression that always returns the same value, such as “5+3” or “ABCDE”
<i>float_expr</i>	Any floating-point expression or an expression that implicitly converts to a floating value
<i>integer_expr</i>	Any integer expression or an expression that implicitly converts to an integer value
<i>numeric_expr</i>	Any numeric expression that returns a single value
<i>char_expr</i>	Any expression that returns a single character-type value
<i>binary_expression</i>	An expression that returns a single binary or varbinary value

## Size of expressions

Expressions returning binary or character datum can be up to 16384 bytes in length. However, earlier versions of Adaptive Server only allowed expressions to be up to 255 bytes in length. If you have upgraded from an earlier release of Adaptive Server, and your stored procedures or scripts store a result string of up to 255 bytes, the remainder will be truncated. You may have to re-write these stored procedures and scripts for to account for the additional length of the expressions.

## Arithmetic and character expressions

The general pattern for arithmetic and character expressions is:

```
{constant | column_name | function | (subquery)
 | (case_expression)}
  [{arithmetic_operator | bitwise_operator |
   string_operator | comparison_operator}
 {constant | column_name | function | (subquery)
 | case_expression}]...
```

## Relational and logical expressions

A logical expression or relational expression returns TRUE, FALSE, or UNKNOWN. The general patterns are:

```
expression comparison_operator [any | all] expression
expression [not] in expression
[not]exists expression
expression [not] between expression and expression
expression [not] like "match_string"
[escape "escape_character "]
not expression like "match_string"
[escape "escape_character "]
expression is [not] null
not logical_expression
logical_expression {and | or} logical_expression
```



## Operator precedence

Operators have the following precedence levels, where 1 is the highest level and 6 is the lowest:

- 1 unary (single argument)  $- + \sim$
- 2  $* / \%$
- 3 binary (two argument)  $+ - \& | ^$
- 4 not
- 5 and
- 6 or

When all operators in an expression are at the same level, the order of execution is left to right. You can change the order of execution with parentheses—the most deeply nested expression is processed first.

## Arithmetic operators

Adaptive Server uses the following arithmetic operators:

**Table 4-2: Arithmetic operators**

Operator	Meaning
+	Addition
-	Subtraction
*	Multiplication
/	Division
%	Modulo (Transact-SQL extension)

Addition, subtraction, division, and multiplication can be used on exact numeric, approximate numeric, and money type columns.

The modulo operator cannot be used on smallmoney, money, numeric, float or real columns. Modulo finds the integer remainder after a division involving two whole numbers. For example,  $21 \% 11 = 10$  because 21 divided by 11 equals 1 with a remainder of 10.

When you perform arithmetic operations on mixed datatypes, for example float and int, Adaptive Server follows specific rules for determining the type of the result. For more information, see Chapter 1, “System and User-Defined Datatypes,”

## Bitwise operators

The bitwise operators are a Transact-SQL extension for use with integer type data. These operators convert each integer operand into its binary representation, then evaluate the operands column by column. A value of 1 corresponds to true; a value of 0 corresponds to false.

Table 4-3 summarizes the results for operands of 0 and 1. If either operand is NULL, the bitwise operator returns NULL:

**Table 4-3: Truth tables for bitwise operations**

& ( and)	1	0
1	1	0
0	0	0
( or)	1	0
1	1	1
0	1	0
^ (exclusive or)	1	0
1	0	1
0	1	0
~ (not)		
1	FALSE	
0	0	

The examples in Table 4-4 use two tinyint arguments, A = 170 (10101010 in binary form) and B = 75 (01001011 in binary form).

**Table 4-4: Examples of bitwise operations**

Operation	Binary form	Result	Explanation
(A & B)	10101010 01001011 ----- 00001010	10	Result column equals 1 if both A and B are 1. Otherwise, result column equals 0.
(A   B)	10101010 01001011 ----- 11101011	235	Result column equals 1 if either A or B, or both, is 1. Otherwise, result column equals 0
(A ^ B)	10101010 01001011 ----- 11100001	225	Result column equals 1 if either A or B, but not both, is 1
(~A)	10101010 ----- 01010101	85	All 1s are changed to 0s and all 0s to 1s

## String concatenation operator

You can use both the + and || (double-pipe) string operators to concatenate two or more character or binary expressions. For example, the following displays author names under the column heading Name in last-name first-name order, with a comma after the last name; for example, “Bennett, Abraham.”:

```
select Name = (au_lname + ", " + au_fname)
from authors
```

This example results in "abcdef", "abcdef":

```
select "abc" + "def", "abc" || "def"
```

The following returns the string “abc def”. The empty string is interpreted as a single space in all char, varchar, unichar, nchar, nvarchar, and text concatenation, and in varchar and univarchar insert and assignment statements:

```
select "abc" + "" + "def"
```

When concatenating non-character, non-binary expressions, always use convert:

```
select "The date is " +
convert(varchar(12), getdate())
```

A string concatenated with NULL evaluates to the value of the string. This is an exception to the SQL standard, which states that a string concatenated with a NULL should evaluate to NULL.

## Comparison operators

Adaptive Server uses the comparison operators listed in Table 4-5:

**Table 4-5: Comparison operators**

Operator	Meaning
=	Equal to
>	Greater than
<	Less than
>=	Greater than or equal to
<=	Less than or equal to
<>	Not equal to
!=	<b>Transact-SQL extension</b> Not equal to
!>	<b>Transact-SQL extension</b> Not greater than
!<	<b>Transact-SQL extension</b> Not less than

In comparing character data, < means closer to the beginning of the server's sort order and > means closer to the end of the sort order. Uppercase and lowercase letters are equal in a case-insensitive sort order. Use `sp_helpsort` to see the sort order for your Adaptive Server. Trailing blanks are ignored for comparison purposes. So, for example, "Dirk" is the same as "Dirk ".

In comparing dates, < means earlier and > means later.

Put single or double quotes around all character and datetime data used with a comparison operator:

```
= "Bennet"  
> "May 22 1947"
```

## Nonstandard operators

The following operators are Transact-SQL extensions:

- Modulo operator: %
- Negative comparison operators: !>, !<, !=

- Bitwise operators: ~, ^, |, &
- Join operators: \*= and =\*

## Using *any*, *all* and *in*

*any* is used with <, >, or = and a subquery. It returns results when any value retrieved in the subquery matches the value in the where or having clause of the outer statement. For more information, see the *Transact-SQL User's Guide*.

*all* is used with < or > and a subquery. It returns results when all values retrieved in the subquery are less than (<) or greater than (>) the value in the where or having clause of the outer statement. For more information, see the *Transact-SQL User's Guide*.

*in* returns results when any value returned by the second expression matches the value in the first expression. The second expression must be a subquery or a list of values enclosed in parentheses. *in* is equivalent to = *any*. For more information, see where clause in *Reference Manual: Commands*.

## Negating and testing

*not* negates the meaning of a keyword or logical expression.

Use *exists*, followed by a subquery, to test for the existence of a particular result.

## Ranges

*between* is the range-start keyword; *and* is the range-end keyword. The following range is inclusive:

```
where column1 between x and y
```

The following range is not inclusive:

```
where column1 > x and column1 < y
```

## Using nulls in expressions

Use *is null* or *is not null* in queries on columns defined to allow null values.

An expression with a bitwise or arithmetic operator evaluates to NULL if any of the operands are null. For example, the following evaluates to NULL if *column1* is NULL:

```
1 + column1
```

## Comparisons that return TRUE

In general, the result of comparing null values is UNKNOWN, since it is not possible to determine whether NULL is equal (or not equal) to a given value or to another NULL. However, the following cases return TRUE when *expression* is any column, variable or literal, or combination of these, which evaluates as NULL:

- *expression* is null
- *expression* = null
- *expression* = @*x*, where @*x* is a variable or parameter containing NULL. This exception facilitates writing stored procedures with null default parameters.
- *expression* != *n*, where *n* is a literal that does not contain NULL, and *expression* evaluates to NULL.

The negative versions of these expressions return TRUE when the expression does not evaluate to NULL:

- *expression* is not null
- *expression* != null
- *expression* != @*x*

---

**Note** The far right side of these exceptions is a literal null, or a variable or parameter containing NULL. If the far right side of the comparison is an expression (such as @*nullvar* + 1), the entire expression evaluates to NULL.

---

Following these rules, null column values do not join with other null column values. Comparing null column values to other null column values in a where clause always returns UNKNOWN for null values, regardless of the comparison operator, and the rows are not included in the results. For example, this query returns no result rows where *column1* contains NULL in both tables (although it may return other rows):

```
select column1
from table1, table2
```

```
where table1.column1 = table2.column1
```

## Difference between FALSE and UNKNOWN

Although neither FALSE nor UNKNOWN returns values, there is an important logical difference between FALSE and UNKNOWN, because the opposite of false (“not false”) is true. For example, “1 = 2” evaluates to false and its opposite, “1 != 2”, evaluates to true. But “not unknown” is still unknown. If null values are included in a comparison, you cannot negate the expression to get the opposite set of rows or the opposite truth value.

## Using “NULL” as a character string

Only columns for which NULL was specified in the create table statement and into which you have explicitly entered NULL (no quotes), or into which no data has been entered, contain null values. Avoid entering the character string “NULL” (with quotes) as data for a character column. It can only lead to confusion. Use “N/A”, “none”, or a similar value instead. When you want to enter the value NULL explicitly, do *not* use single or double quotes.

## NULL compared to the empty string

The empty string (“ ” or ‘ ’) is always stored as a single space in variables and column data. This concatenation statement is equivalent to “abc def”, not to “abcdef”:

```
"abc" + " " + "def"
```

The empty string is never evaluated as NULL.

## Connecting expressions

and connects two expressions and returns results when both are true. or connects two or more conditions and returns results when either of the conditions is true.

When more than one logical operator is used in a statement, and is evaluated before or. You can change the order of execution with parentheses.

Table 4-6 shows the results of logical operations, including those that involve null values.

**Table 4-6: Truth tables for logical expressions**

and	TRUE	FALSE	NULL
TRUE	TRUE	FALSE	UNKNOWN
FALSE	FALSE	FALSE	FALSE
NULL	UNKNOWN	FALSE	UNKNOWN
or	TRUE	FALSE	NULL
TRUE	TRUE	TRUE	TRUE
FALSE	TRUE	FALSE	UNKNOWN
NULL	TRUE	UNKNOWN	UNKNOWN
not			
TRUE	FALSE		
FALSE	TRUE		
NULL	UNKNOWN		

The result UNKNOWN indicates that one or more of the expressions evaluates to NULL, and that the result of the operation cannot be determined to be either TRUE or FALSE. See “Using nulls in expressions” on page 281 for more information.

## Using parentheses in expressions

Parentheses can be used to group the elements in an expression. When “expression” is given as a variable in a syntax statement, a simple expression is assumed. “Logical expression” is specified when only a logical expression is acceptable.

## Comparing character expressions

Character constant expressions are treated as varchar. If they are compared with non-varchar variables or column data, the datatype precedence rules are used in the comparison (that is, the datatype with lower precedence is converted to the datatype with higher precedence). If implicit datatype conversion is not supported, you must use the convert function.

Comparison of a char expression to a varchar expression follows the datatype precedence rule; the “lower” datatype is converted to the “higher” datatype. All varchar expressions are converted to char (that is, trailing blanks are appended) for the comparison. If a unichar expression is compared to a char (varchar, nchar, nvarchar) expression, the latter is implicitly converted to unichar.



## Using the empty string

The empty string ("") or (' ') is interpreted as a single blank in insert or assignment statements on varchar or univarchar data. In concatenation of varchar, char, nchar, nvarchar data, the empty string is interpreted as a single space; for following example is stored as "abc def":

```
"abc" + "" + "def"
```

The empty string is never evaluated as NULL.

## Including quotation marks in character expressions

There are two ways to specify literal quotes within a char, or varchar entry. The first method is to double the quotes. For example, if you begin a character entry with a single quote and you want to include a single quote as part of the entry, use two single quotes:

```
'I don''t understand.'
```

With double quotes:

```
"He said, ""It's not really confusing."""
```

The second method is to enclose a quote in the opposite kind of quote mark. In other words, surround an entry containing a double quote with single quotes (or vice versa). Here are some examples:

```
'George said, "There must be a better way."'
'Isn't there a better way?'
'George asked, "Isn't there a better way?'"
```

## Using the continuation character

To continue a character string to the next line on your screen, enter a backslash (\) before going to the next line.

## Identifiers

Identifiers are names for database objects such as databases, tables, views, columns, indexes, triggers, procedures, defaults, rules, and cursors.

The limit for the length of object names or identifiers is 255 bytes for regular identifiers, and 253 bytes for delimited identifiers. The limit applies to most user-defined identifiers including table name, column name, index name and so on. Due to the expanded limits, some system tables (catalogs) and built-in functions have been expanded.

For variables, “@” count as 1 byte, and the allowed name for it is 254 bytes long.

Listed below are the identifiers, system tables, and built-in functions that are affected these limits.

The maximum length for these identifiers is now 255 bytes.

- Table name
- Column name
- Index name
- View name
- User-defined datatype
- Trigger name
- Default name
- Rule name
- Constraint name
- Procedure name
- Variable name
- JAR name
- Name of LWP or dynamic statement
- Function name
- Name of the time range
- Application context name

Most user-defined Adaptive Server identifiers can be a maximum of 255 bytes in length, whether single-byte or multibyte characters are used. Others can be a maximum of 30 bytes. Refer to the *Transact-SQL User's Guide* for a list of both 255-byte and 30-byte identifiers.

The first character of an identifier must be either an alphabetic character, as defined in the current character set, or the underscore (`_`) character.

---

**Note** Temporary table names, which begin with the pound sign (`#`), and variable names, which begin with the at sign (`@`), are exceptions to this rule.

---

Subsequent characters can include letters, numbers, the symbols `#`, `@`, `_`, and currency symbols such as `$` (dollars), `¥` (yen), and `£` (pound sterling). Identifiers cannot include special characters such as `!`, `%`, `^`, `&`, `*`, and `.` or embedded spaces.

You cannot use a reserved word, such as a Transact-SQL command, as an identifier. For a complete list of reserved words, see Chapter 5, “Reserved Words.”

You cannot use the dash symbol (`-`) as an identifier.

## Short identifiers

The maximum length for these identifiers is 30 bytes:

- Cursor name
- Server name
- Host name
- Login name
- Password
- Host process identification
- Application name
- Initial language name
- Character set name
- User name
- Group name
- Database name
- Logical device name
- Segment name

- Session name
- Execution class name
- Engine name
- Quiesce tag name
- Cache name

## Tables beginning with # (temporary tables)

Tables with names that begin with the pound sign (#) are temporary tables. You cannot create other types of objects with names that begin with the pound sign.

Adaptive Server performs special operations on temporary table names to maintain unique naming on a per-session basis. When you create a temporary table with a name of fewer than 238 bytes, the sysobjects name in the tempdb adds 17 bytes to make the table name unique. If the table name is more than 238 bytes, the temporary table name in sysobjects uses only the first 238 bytes, then adds 17 bytes to make it unique.

In versions of Adaptive Server earlier than 15.0, temporary table names in sysobjects were 30 bytes. If you used a table name with fewer than 13 bytes, the name was padded with underscores ( \_ ) to 13 bytes, then another 17 bytes of other characters to bring the name up to 30 bytes.

## Case sensitivity and identifiers

Sensitivity to the case (upper or lower) of identifiers and data depends on the sort order installed on your Adaptive Server. Case sensitivity can be changed for single-byte character sets by reconfiguring Adaptive Server's sort order; see the *System Administration Guide* for more information. Case is significant in utility program options.

If Adaptive Server is installed with a case-insensitive sort order, you cannot create a table named MYTABLE if a table named MyTable or mytable already exists. Similarly, the following command will return rows from MYTABLE, MyTable, or mytable, or any combination of uppercase and lowercase letters in the name:

```
select * from MYTABLE
```

## Uniqueness of object names

Object names need not be unique in a database. However, column names and index names must be unique within a table, and other object names must be unique for each *owner* within a *database*. Database names must be unique on Adaptive Server.

## Using delimited identifiers

**Delimited identifiers** are object names enclosed in double quotes. Using delimited identifiers allows you to avoid certain restrictions on object names. Table, view, and column names can be delimited by quotes; other object names cannot.

Delimited identifiers can be reserved words, can begin with non-alphabetic characters, and can include characters that would not otherwise be allowed. They cannot exceed 253 bytes.

---

**Warning!** Delimited identifiers may not be recognized by all front-end applications and should not be used as parameters to system procedures.

---

Before creating or referencing a delimited identifier, you must execute:

```
set quoted_identifier on
```

Each time you use the delimited identifier in a statement, you must enclose it in double quotes. For example:

```
create table "1one" (col1 char(3))
create table "include spaces" (col1 int)
create table "grant" ("add" int)
insert "grant" ("add") values (3)
```

While the `quoted_identifier` option is turned on, do not use double quotes around character or date strings; use single quotes instead. Delimiting these strings with double quotes causes Adaptive Server to treat them as identifiers. For example, to insert a character string into *col1* of *1table*, use:

```
insert "1one" (col1) values ('abc')
```

Do not use:

```
insert "1one" (col1) values ("abc")
```

To insert a single quote into a column, use two consecutive single quotation marks. For example, to insert the characters “a'b” into *col1* use:

Syntax that includes quotes

```
insert "lone" (col1) values ('a' 'b')
```

When the `quoted_identifier` option is set to on, you do not need to use double quotes around an identifier if the syntax of the statement requires that a quoted string contain an identifier. For example:

```
set quoted_identifier on
create table 'lone' (c1 int)
```

However, `object_id()` requires a string, so you must include the table name in quotes to select the information:

```
select object_id('lone')
-----
      896003192
```

You can include an embedded double quote in a quoted identifier by doubling the quote:

```
create table "embedded" "quote" (c1 int)
```

However, there is no need to double the quote when the statement syntax requires the object name to be expressed as a string:

```
select object_id('embedded"quote')
```

## Identifying tables or columns by their qualified object name

You can uniquely identify a table or column by adding other names that qualify it—the database name, owner's name, and (for a column) the table or view name. Each qualifier is separated from the next one by a period. For example:

```
database.owner.table_name.column_name
database.owner.view_name.column_name
```

The naming conventions are:

```
[[database.]owner.]table_name
[[database.]owner.]view_name
```

## Using delimited identifiers within an object name

If you use `set quoted_identifier on`, you can use double quotes around individual parts of a qualified object name. Use a separate pair of quotes for each qualifier that requires quotes. For example, use:

```
database.owner."table_name"."column_name"
```

Do not use:

```
database.owner."table_name.column_name"
```

### Omitting the owner name

You can omit the intermediate elements in a name and use dots to indicate their positions, as long as the system is given enough information to identify the object:

```
database..table_name
```

```
database..view_name
```

### Referencing your own objects in the current database

You need not use the database name or owner name to reference your own objects in the current database. The default value for *owner* is the current user, and the default value for *database* is the current database.

If you reference an object without qualifying it with the database name and owner name, Adaptive Server tries to find the object in the current database among the objects you own.

### Referencing objects owned by the database owner

If you omit the owner name and you do not own an object by that name, Adaptive Server looks for objects of that name owned by the Database Owner. You must qualify objects owned by the Database Owner only if you own an object of the same name, but you want to use the object owned by the Database Owner. However, you must qualify objects owned by other users with the user's name, whether or not you own objects of the same name.

### Using qualified identifiers consistently

When qualifying a column name and table name in the same statement, be sure to use the same qualifying expressions for each; they are evaluated as strings and must match; otherwise, an error is returned. Example 2 is incorrect because the syntax style for the column name does not match the syntax style used for the table name.

Example 1

```
select demo.mary.publishers.city  
from demo.mary.publishers
```

```
city
-----
Boston
Washington
Berkeley
```

Example 2

```
select demo.mary.publishers.city
from demo..publishers
```

The column prefix "demo.mary.publishers" does not match a table name or alias name used in the query.

## Determining whether an identifier is valid

Use the system function `valid_name`, after changing character sets or before creating a table or view, to determine whether the object name is acceptable to Adaptive Server. Here is the syntax:

```
select valid_name("Object_name")
```

If *object\_name* is not a valid identifier (for example, if it contains illegal characters or is more than 30 bytes long), Adaptive Server returns 0. If *object\_name* is a valid identifier, Adaptive Server returns a nonzero number.

## Renaming database objects

Rename user objects (including user-defined datatypes) with `sp_rename`.

---

**Warning!** After you rename a table or column, you must redefine all procedures, triggers, and views that depend on the renamed object.

---

## Using multibyte character sets

In multibyte character sets, a wider range of characters is available for use in identifiers. For example, on a server with the Japanese language installed, the following types of characters may be used as the first character of an identifier: Zenkaku or Hankaku Katakana, Hiragana, Kanji, Romaji, Greek, Cyrillic, or ASCII.



Although Hankaku Katakana characters are legal in identifiers on Japanese systems, they are not recommended for use in heterogeneous systems. These characters cannot be converted between the EUC-JIS and Shift-JIS character sets.

The same is true for some 8-bit European characters. For example, the OE ligature, is part of the Macintosh character set (codepoint 0xCE). This character does not exist in the ISO 8859-1 (iso\_1) character set. If the OE ligature exists in data being converted from the Macintosh to the ISO 8859-1 character set, it causes a conversion error.

If an object identifier contains a character that cannot be converted, the client loses direct access to that object.

## Pattern matching with wildcard characters

Wildcard characters represent one or more characters, or a range of characters, in a *match\_string*. A *match\_string* is a character string containing the pattern to find in the expression. It can be any combination of constants, variables, and column names or a concatenated expression, such as:

```
like @variable + "%".
```

If the match string is a constant, it must always be enclosed in single or double quotes.

Use wildcard characters with the keyword `like` to find character and date strings that match a particular pattern. You cannot use `like` to search for seconds or milliseconds. For more information, see “Using wildcard characters with datetime data” on page 299.

Use wildcard characters in `where` and `having` clauses to find character or date/time information that is `like`—or not `like`—the match string:

```
{where | having} [not]
  expression [not] like match_string
  [escape "escape_character"]
```

*expression* can be any combination of column names, constants, or functions with a character value.

Wildcard characters used without `like` have no special meaning. For example, this query finds any phone numbers that start with the four characters “415%”:

```
select phone
```

```
from authors
where phone = "415%"
```

## Using *not like*

Use *not like* to find strings that do not match a particular pattern. These two queries are equivalent: they find all the phone numbers in the authors table that do not begin with the 415 area code.

```
select phone
from authors
where phone not like "415%"
```

```
select phone
from authors
where not phone like "415%"
```

For example, this query finds the system tables in a database whose names begin with “sys”:

```
select name
from sysobjects
where name like "sys%"
```

To see all the objects that are *not* system tables, use:

```
not like "sys%"
```

If you have a total of 32 objects and *like* finds 13 names that match the pattern, *not like* will find the 19 objects that do not match the pattern.

*not like* and the negative wildcard character [^] may give different results (see “The caret (^) wildcard character” on page 297). You cannot always duplicate *not like* patterns with *like* and ^. This is because *not like* finds the items that do not match the entire *like* pattern, but *like* with negative wildcard characters is evaluated one character at a time.

A pattern such as *like* “[^s][^y][^s]%" may not produce the same results. Instead of 19, you might get only 14, with all the names that begin with “s”, *or* have “y” as the second letter, *or* have “s” as the third letter eliminated from the results, as well as the system table names. This is because match strings with negative wildcard characters are evaluated in steps, one character at a time. If the match fails at any point in the evaluation, it is eliminated.

## Case and accent insensitivity

If your Adaptive Server uses a case-insensitive sort order, case is ignored when comparing *expression* and *match\_string*. For example, this clause would return “Smith,” “smith,” and “SMITH” on a case-insensitive Adaptive Server:

```
where col_name like "Sm%"
```

If your Adaptive Server is also accent-insensitive, it treats all accented characters as equal to each other and to their unaccented counterparts, both uppercase and lowercase. The `sp_helpsort` system procedure displays the characters that are treated as equivalent, displaying an “=” between them.

## Using wildcard characters

You can use the match string with a number of wildcard characters, which are discussed in detail in the following sections. Table 4-7 summarizes the wildcard characters:

**Table 4-7: Wildcard characters used with like**

Symbol	Meaning
%	Any string of 0 or more characters
_	Any single character
[ ]	Any single character within the specified range ([a-f]) or set ([abcdef])
[^]	Any single character not within the specified range ([^a-f]) or set ([^abcdef])

Enclose the wildcard character and the match string in single or double quotes (like “[dD]eFr\_nce”).

### The percent sign (%) wildcard character

Use the % wildcard character to represent any string of zero or more characters. For example, to find all the phone numbers in the authors table that begin with the 415 area code:

```
select phone
from authors
where phone like "415%"
```

To find names that have the characters “en” in them (Bennet, Green, McBaden):

```
select au_lname
from authors
```

```
where au_lname like "%en%"
```

Trailing blanks following “%” in a like clause are truncated to a single trailing blank. For example, “%” followed by two spaces matches “X ”(one space); “X ” (two spaces); “X ” (three spaces), or any number of trailing spaces.

## The underscore ( \_ ) wildcard character

Use the underscore ( \_ ) wildcard character to represent any single character. For example, to find all six-letter names that end with “heryl” (for example, Cheryl):

```
select au_fname
from authors
where au_fname like "_heryl"
```

## Bracketed ( [ ] ) characters

Use brackets to enclose a range of characters, such as [a-f], or a set of characters such as [a2Br]. When ranges are used, all values in the sort order between (and including) *rangespec1* and *rangespec2* are returned. For example, “[0-z]” matches 0-9, A-Z and a-z (and several punctuation characters) in 7-bit ASCII.

To find names ending with “inger” and beginning with any single character between M and Z:

```
select au_lname
from authors
where au_lname like "[M-Z]inger"
```

To find both “DeFrance” and “deFrance”:

```
select au_lname
from authors
where au_lname like "[dD]eFrance"
```

When using bracketed identifiers to create objects, such as with `create table [table_name]` or `create database [dbname]`, you must include at least one valid character.

All trailing spaces within bracketed identifiers are removed from the object name. For example, you achieve the same results executing the following `create table` commands:

- `create table [tab1<space><space>]`
- `create table [tab1]`

- create table [tab1<space><space><space>]
- create table tab1

This rule applies to all objects you can create using bracketed identifiers.

## The caret (^) wildcard character

The caret is the negative wildcard character. Use it to find strings that do not match a particular pattern. For example, “[^a-f]” finds strings that are not in the range a-f and “[^a2bR]” finds strings that are not “a,” “2,” “b,” or “R.”

To find names beginning with “M” where the second letter is not “c”:

```
select au_lname
from authors
where au_lname like "M[^c]%"
```

When ranges are used, all values in the sort order between (and including) *rangespec1* and *rangespec2* are returned. For example, “[0-z]” matches 0-9, A-Z, a-z, and several punctuation characters in 7-bit ASCII.

## Using multibyte wildcard characters

If the multibyte character set configured on your Adaptive Server defines equivalent double-byte characters for the wildcard characters `_`, `%`, `- [`, `]`, and `^`, you can substitute the equivalent character in the match string. The underscore equivalent represents either a single- or double-byte character in the match string.

## Using wildcard characters as literal characters

To search for the occurrence of `%`, `_`, `[`, `]`, or `^` within a string, you must use an escape character. When a wildcard character is used in conjunction with an escape character, Adaptive Server interprets the wildcard character literally, rather than using it to represent other characters.

Adaptive Server provides two types of escape characters:

- Square brackets, a Transact-SQL extension
- Any single character that immediately follows an escape clause, compliant with the SQL standards

## Using square brackets ( [ ] ) as escape characters

Use square brackets as escape characters for the percent sign, the underscore, and the left bracket. The right bracket does not need an escape character; use it by itself. If you use the hyphen as a literal character, it must be the first character inside a set of square brackets.

Table 4-8 shows examples of square brackets used as escape characters with like.

**Table 4-8: Using square brackets to search for wildcard characters**

like predicate	Meaning
like "5%"	5 followed by any string of 0 or more characters
like "5[%]"	5%
like "_n"	an, in, on (and so on)
like "[_n]"	_n
like "[a-cdf]"	a, b, c, d, or f
like "[-acdf]"	-, a, c, d, or f
like "[["	[
like "]"	]
like "[[ab]"	[ab

## Using the escape clause

Use the escape clause to specify an escape character. Any single character in the server's default character set can be used as an escape character. If you try to use more than one character as an escape character, Adaptive Server generates an exception.

Do not use existing wildcard characters as escape characters because:

- If you specify the underscore ( \_ ) or percent sign ( % ) as an escape character, it loses its special meaning within that like predicate and acts only as an escape character.
- If you specify the left or right bracket ( [ or ] ) as an escape character, the Transact-SQL meaning of the bracket is disabled within that like predicate.
- If you specify the hyphen ( - ) or caret ( ^ ) as an escape character, it loses its special meaning and acts only as an escape character.

An escape character retains its special meaning within square brackets, unlike wildcard characters such as the underscore, the percent sign, and the open bracket.

The escape character is valid only within its like predicate and has no effect on other like predicates contained in the same statement. The only characters that are valid following an escape character are the wildcard characters ( `_`, `%`, `[`, `]`, or `[^]` ), and the escape character itself. The escape character affects only the character following it, and subsequent characters are not affected by it.

If the pattern contains two literal occurrences of the character that happens to be the escape character, the string must contain four consecutive escape characters. If the escape character does not divide the pattern into pieces of one or two characters, Adaptive Server returns an error message. Table 4-9 shows examples of escape clauses used with like.

**Table 4-9: Using the escape clause**

like predicate	Meaning
like "5@%" escape "@"	5%
like "*_n" escape "**"	_n
like "%80@%" escape "@"	String containing 80%
like "*_sql**%" escape "**"	String containing _sql*
like "%#####_#%" escape "#"	String containing ##_%

## Using wildcard characters with *datetime* data

When you use like with datetime values, Adaptive Server converts the dates to the standard datetime format, then to varchar. Since the standard storage format does not include seconds or milliseconds, you cannot search for seconds or milliseconds with like and a pattern.

It is a good idea to use like when you search for datetime values, since datetime entries may contain a variety of date parts. For example, if you insert the value “9:20” and the current date into a column named arrival\_time, the clause:

```
where arrival_time = '9:20'
```

would not find the value, because Adaptive Server converts the entry into “Jan 1 1900 9:20AM.” However, the following clause would find this value:

```
where arrival_time like '%9:20%'
```





Keywords, also known as reserved words, are words that have special meanings. This chapter lists Transact-SQL and ANSI SQL keywords.

Topics covered are:

Topics	Page
Transact-SQL reserved words	301
ANSI SQL reserved words	302
Potential ANSI SQL reserved words	303

## Transact-SQL reserved words

The words in Table 5-1 are reserved by Adaptive Server as keywords (part of SQL command syntax). They cannot be used as names of database objects such as databases, tables, rules, or defaults. They can be used as names of local variables and as stored procedure parameter names.

To find the names of existing objects that are reserved words, use `sp_checkreswords` in *Reference Manual: Procedures*.

**Table 5-1: List of Transact-SQL reserved words**

	Words
<i>A</i>	add, all, alter, and, any, arith_overflow, as, asc, at, authorization, avg
<i>B</i>	begin, between, break, browse, bulk, by
<i>C</i>	cascade, case, char_convert, check, checkpoint, close, clustered, coalesce, commit, compute, confirm, connect, constraint, continue, controlrow, convert, count, count_big, create, current, cursor
<i>D</i>	database, dbcc, deallocate, declare, decrypt, default, delete, desc, deterministic, disk, distinct, drop, dummy, dump
<i>E</i>	else, encrypt, end, endtran, errlvl, errordata, errexit, escape, except, exclusive, exec, execute, exists, exit, exp_row_size, external
<i>F</i>	fetch, fillfactor, for, foreign, from
<i>G</i>	goto, grant, group
<i>H</i>	having, holdlock

	Words
<i>I</i>	identity, identity_gap, identity_start, if, in, index, inout, insensitive, insert, install, intersect, into, is, isolation
<i>J</i>	jar, join
<i>K</i>	key, kill
<i>L</i>	level, like, lineno, load, lock
<i>M</i>	materialized, max, max_rows_per_page, min, mirror, mirrorexit, modify
<i>N</i>	national, new, noholdlock, nonclustered, nonscrollable, non_sensitive, not, null, nullif, numeric_truncation  <p>Note Although “new” is not a Transact-SQL reserved word, since it may become a reserved word in the future, Sybase recommends that you avoid using it (for example, to name a database object). “New” is a special case (see “Potential ANSI SQL reserved words” on page 303 for information on other reserved words) because it appears in the <code>spt_values</code> table, and because <code>sp_checkreswords</code> displays “New” as a reserved word.</p>
<i>O</i>	of, off, offsets, on, once, online, only, open, option, or, order, out, output, over
<i>P</i>	partition, perm, permanent, plan, prepare, primary, print, privileges, proc, procedure, processexit, proxy_table, public
<i>Q</i>	quiesce
<i>R</i>	raiserror, read, readpast, readtext, reconfigure, references, remove, reorg, replace, replication, reservepagegap, return, returns, revoke, role, rollback, rowcount, rows, rule
<i>S</i>	save, schema, scroll, scrollable, select, semi_sensitive, set, setuser, shared, shutdown, some, statistics, stringsize, stripe, sum, syb_identity, syb_restree, syb_terminate
<i>T</i>	table, temp, temporary, textsize, to, tracefile, tran, transaction, trigger, truncate, tsequal
<i>U</i>	union, unique, unpartition, update, use, user, user_option, using
<i>V</i>	values, varying, view
<i>W</i>	waitfor, when, where, while, with, work, writetext
<i>X</i>	xmlextract, xmlparse, xmltest, xmlvalidate

## ANSI SQL reserved words

Adaptive Server includes entry-level ANSI SQL features. Full ANSI SQL implementation includes the words listed in the following tables as command syntax. Upgrading identifiers can be a complex process; therefore, we are providing this list for your convenience. The publication of this information does not commit Sybase to providing all of these ANSI SQL features in subsequent releases. In addition, subsequent releases may include keywords not included in this list.

The words in Table 5-2 are ANSI SQL keywords that are not reserved words in Transact-SQL.

**Table 5-2: List of ANSI SQL reserved words**

<b>Words</b>	
<i>A</i>	absolute, action, allocate, are, assertion
<i>B</i>	bit, bit_length, both
<i>C</i>	cascaded, case, cast, catalog, char, char_length, character, character_length, coalesce, collate, collation, column, connection, constraints, corresponding, cross, current_date, current_time, current_timestamp, current_user
<i>D</i>	date, day, dec, decimal, deferrable, deferred, describe, descriptor, diagnostics, disconnect, domain
<i>E</i>	end-exec, exception, extract
<i>F</i>	false, first, float, found, full
<i>G</i>	get, global, go
<i>H</i>	hour
<i>I</i>	immediate, indicator, initially, inner, input, insensitive, int, integer, interval
<i>J</i>	join
<i>L</i>	language, last, leading, left, local, lower
<i>M</i>	match, minute, module, month
<i>N</i>	names, natural, nchar, next, no, nullif, numeric
<i>O</i>	octet_length, outer, output, overlaps
<i>P</i>	pad, partial, position, preserve, prior
<i>R</i>	real, relative, restrict, right
<i>S</i>	scroll, second, section, semi_sensitive, session_user, size, smallint, space, sql, sqlcode, sqlerror, sqlstate, substring, system_user
<i>T</i>	then, time, timestamp, timezone_hour, timezone_minute, trailing, translate, translation, trim, true
<i>U</i>	unknown, upper, usage
<i>V</i>	value, varchar
<i>W</i>	when, whenever, write, year
<i>Z</i>	zone

## Potential ANSI SQL reserved words

If you are using the ISO/IEC 9075:1989 standard, also avoid using the words shown in the following list because these words may become ANSI SQL reserved words in the future.

**Table 5-3: List of potential ANSI SQL reserved words**

	<b>Words</b>
<i>A</i>	after, alias, async
<i>B</i>	before, boolean, breadth
<i>C</i>	call, completion, cycle
<i>D</i>	data, depth, dictionary
<i>E</i>	each, elseif, equals
<i>G</i>	general
<i>I</i>	ignore
<i>L</i>	leave, less, limit, loop
<i>M</i>	modify
<i>N</i>	new, none
<i>O</i>	object, oid, old, operation, operators, others
<i>P</i>	parameters, pendant, preorder, private, protected
<i>R</i>	recursive, ref, referencing, resignal, return, returns, routine, row
<i>S</i>	savepoint, search, sensitive, sequence, signal, similar, sqlexception, structure
<i>T</i>	test, there, type
<i>U</i>	under
<i>V</i>	variable, virtual, visible
<i>W</i>	wait, without

# SQLSTATE Codes and Messages

This chapter describes Adaptive Server's SQLSTATE status codes and their associated messages.

Topics covered are:

Topics	Page
Warnings	305
Exceptions	306

SQLSTATE codes are required for entry level ANSI SQL compliance. They provide diagnostic information about two types of conditions:

- *Warnings* – conditions that require user notification but are not serious enough to prevent a SQL statement from executing successfully
- *Exceptions* – conditions that prevent a SQL statement from having any effect on the database

Each SQLSTATE code consists of a 2-character class followed by a 3-character subclass. The class specifies general information about error type. The subclass specifies more specific information.

SQLSTATE codes are stored in the sysmessages system table, along with the messages that display when these conditions are detected. Not all Adaptive Server error conditions are associated with a SQLSTATE code—only those mandated by ANSI SQL. In some cases, multiple Adaptive Server error conditions are associated with a single SQLSTATE value.

## Warnings

Adaptive Server currently detects the following SQLSTATE warning conditions, described in Table 6-1:

**Table 6-1: SQLSTATE warnings**

Message	Value	Description
Warning – null value eliminated in set function.	01003	Occurs when you use an aggregate function (avg, max, min, sum, or count) on an expression with a null value.
Warning–string data, right truncation	01004	Occurs when character, unichar, or binary data is truncated to 255 bytes. The data may be: <ul style="list-style-type: none"> <li>• The result of a select statement in which the client does not support the WIDE TABLES property.</li> <li>• Parameters to an RPC on remote Adaptive Servers or Open Servers that do not support the WIDE TABLES property.</li> </ul>

## Exceptions

Adaptive Server detects the following types of exceptions:

- Cardinality violations
- Data exceptions
- Integrity constraint violations
- Invalid cursor states
- Syntax errors and access rule violations
- Transaction rollbacks
- with check option violations

Exception conditions are described in Table 6-2 through Table 6-8. Each class of exceptions appears in its own table. Within each table, conditions are sorted alphabetically by message text.

### Cardinality violations

Cardinality violations occur when a query that should return only a single row returns more than one row to an Embedded SQL™ application.

**Table 6-2: Cardinality violations**

Message	Value	Description
Subquery returned more than 1 value. This is illegal when the subquery follows =, !=, <, <=, >, >=. or when the subquery is used as an expression.	21000	Occurs when: <ul style="list-style-type: none"> <li>• A scalar subquery or a row subquery returns more than one row.</li> <li>• A select into parameter_list query in Embedded SQL returns more than one row.</li> </ul>

## Data exceptions

Data exceptions occur when an entry:

- Is too long for its datatype,
- Contains an illegal escape sequence, or
- Contains other format errors.

**Table 6-3: Data exceptions**

Message	Value	Description
Arithmetic overflow occurred.	22003	Occurs when: <ul style="list-style-type: none"> <li>• An exact numeric type would lose precision or scale as a result of an arithmetic operation or sum function.</li> <li>• An approximate numeric type would lose precision or scale as a result of truncation, rounding, or a sum function.</li> </ul>
Data exception - string data right truncated.	22001	Occurs when a char, unichar, univarchar, or varchar column is too short for the data being inserted or updated and non-blank characters must be truncated.
Divide by zero occurred.	22012	Occurs when a numeric expression is being evaluated and the value of the divisor is zero.
Illegal escape character found. There are fewer bytes than necessary to form a valid character.	22019	Occurs when you are searching for strings that match a given pattern if the escape sequence does not consist of a single character.
Invalid pattern string. The character following the escape character must be percent sign, underscore, left square bracket, right square bracket, or the escape character.	22025	Occurs when you are searching for strings that match a particular pattern when: <ul style="list-style-type: none"> <li>• The escape character is not immediately followed by a percent sign, an underscore, or the escape character itself, or</li> <li>• The escape character partitions the pattern into substrings whose lengths are other than 1 or 2 characters.</li> </ul>

## Integrity constraint violations

Integrity constraint violations occur when an insert, update, or delete statement violates a primary key, foreign key, check, or unique constraint or a unique index.

**Table 6-4: Integrity constraint violations**

Message	Value	Description
Attempt to insert duplicate key row in object <i>object_name</i> with unique index <i>index_name</i> .	23000	Occurs when a duplicate row is inserted into a table that has a unique constraint or index.
Check constraint violation occurred, dbname = <i>database_name</i> , table name = <i>table_name</i> , constraint name = <i>constraint_name</i> .	23000	Occurs when an update or delete would violate a check constraint on a column.
Dependent foreign key constraint violation in a referential integrity constraint. dbname = <i>database_name</i> , table name = <i>table_name</i> , constraint name = <i>constraint_name</i> .	23000	Occurs when an update or delete on a primary key table would violate a foreign key constraint.
Foreign key constraint violation occurred, dbname = <i>database_name</i> , table name = <i>table_name</i> , constraint name = <i>constraint_name</i> .	23000	Occurs when an insert or update on a foreign key table is performed without a matching value in the primary key table.

## Invalid cursor states

Invalid cursor states occur when:

- A fetch uses a cursor that is not currently open, or
- An update where current of or delete where current of affects a cursor row that has been modified or deleted, or
- An update where current of or delete where current of affects a cursor row that not been fetched.

**Table 6-5: Invalid cursor states**

Message	Value	Description
Attempt to use cursor <i>cursor_name</i> which is not open. Use the system stored procedure <i>sp_cursorinfo</i> for more information.	24000	Occurs when an attempt is made to fetch from a cursor that has never been opened or that was closed by a commit statement or an implicit or explicit rollback. Reopen the cursor and repeat the fetch.



Message	Value	Description
Cursor <i>cursor_name</i> was closed implicitly because the current cursor position was deleted due to an update or a delete. The cursor scan position could not be recovered. This happens for cursors which reference more than one table.	24000	Occurs when the join column of a multitable cursor has been deleted or changed. Issue another fetch to reposition the cursor.
The cursor <i>cursor_name</i> had its current scan position deleted because of a DELETE/UPDATE WHERE CURRENT OF or a regular searched DELETE/UPDATE. You must do a new FETCH before doing an UPDATE or DELETE WHERE CURRENT OF.	24000	Occurs when a user issues an update/delete where current of whose current cursor position has been deleted or changed. Issue another fetch before retrying the update/delete where current of.
The UPDATE/DELETE WHERE CURRENT OF failed for the cursor <i>cursor_name</i> because it is not positioned on a row.	24000	Occurs when a user issues an update/delete where current of on a cursor that: <ul style="list-style-type: none"> <li>• Has not yet fetched a row</li> <li>• Has fetched one or more rows after reaching the end of the result set</li> </ul>

## Syntax errors and access rule violations

Syntax errors are generated by SQL statements that contain unterminated comments, implicit datatype conversions not supported by Adaptive Server or other incorrect syntax.

Access rule violations are generated when a user tries to access an object that does not exist or one for which he or she does not have the correct permissions.

**Table 6-6: Syntax errors and access rule violations**

Message	Value	Description
<i>command</i> permission denied on object <i>object_name</i> , database <i>database_name</i> , owner <i>owner_name</i> .	42000	Occurs when a user tries to access an object for which he or she does not have the proper permissions.
Implicit conversion from datatype ' <i>datatype</i> ' to ' <i>datatype</i> ' is not allowed. Use the CONVERT function to run this query.	42000	Occurs when the user attempts to convert one datatype to another but Adaptive Server cannot do the conversion implicitly.
Incorrect syntax near <i>object_name</i> .	42000	Occurs when incorrect SQL syntax is found near the object specified.

Message	Value	Description
Insert error: column name or number of supplied values does not match table definition.	42000	Occurs during inserts when an invalid column name is used or when an incorrect number of values is inserted.
Missing end comment mark `*/'.	42000	Occurs when a comment that begins with the /* opening delimiter does not also have the */ closing delimiter.
<i>object_name</i> not found. Specify <i>owner.objectname</i> or use <i>sp_help</i> to check whether the object exists ( <i>sp_help</i> may produce lots of output).	42000	Occurs when a user tries to reference an object that he or she does not own. When referencing an object owned by another user, be sure to qualify the object name with the name of its owner.
The size ( <i>size</i> ) given to the <i>object_name</i> exceeds the maximum. The largest size allowed is <i>size</i> .	42000	Occurs when: <ul style="list-style-type: none"> <li>• The total size of all the columns in a table definition exceeds the maximum allowed row size.</li> <li>• The size of a single column or parameter exceeds the maximum allowed for its datatype.</li> </ul>

## Transaction rollbacks

Transaction rollbacks occur when the transaction isolation level is set to 3, but Adaptive Server cannot guarantee that concurrent transactions can be serialized. This type of exception generally results from system problems such as disk crashes and offline disks.

**Table 6-7: Transaction rollbacks**

Message	Value	Description
Your server command (process id # <i>process_id</i> ) was deadlocked with another process and has been chosen as deadlock victim. Re-run your command.	40001	Occurs when Adaptive Server detects that it cannot guarantee that two or more concurrent transactions can be serialized.

### *with check option* violation

This class of exception occurs when data being inserted or updated through a view would not be visible through the view.

**Table 6-8: with check option violation**

<b>Message</b>	<b>Value</b>	<b>Description</b>
The attempted insert or update failed because the target view was either created WITH CHECK OPTION or spans another view created WITH CHECK OPTION. At least one resultant row from the command would not qualify under the CHECK OPTION constraint.	44000	Occurs when a view, or any view on which it depends, was created with a with check option clause.



# Index

## Symbols

- & (ampersand) “and” bitwise operator 278
- \* (asterisk)
  - for overlength numbers 232
  - multiplication operator 277
- \ (backslash) character string continuation with 285
- ::= (BNF notation)
  - in SQL statements xvii
- ^ (caret)
  - “exclusive or” bitwise operator 278
  - wildcard character 295, 297
- : (colon) preceding milliseconds 65, 133
- , (comma)
  - in default print format for money values 18
  - not allowed in money values 19
  - in SQL statements xvii
- { } (curly braces)
  - in SQL statements xvii
- \$ (dollar sign)
  - in identifiers 287
  - in money datatypes 19
- .. (dots) in database object names 291
- || (double pipe)
  - string concatenation operator 279
- = (equals sign) comparison operator 280
- > (greater than) comparison operator 280
- >= (greater than or equal to) comparison operator 280
- < (less than) comparison operator 280
- <= (less than or equal to) comparison operator 280
- (minus sign)
  - arithmetic operator 277
  - for negative monetary values 19
  - in integer data 13
- != (not equal to) comparison operator 280
- <> (not equal to) comparison operator 280
- !> (not greater than) comparison operator 280
- !< (not less than) comparison operator 280
- () (parentheses)
  - in expressions 284
  - in SQL statements xvii
- % (percent sign)
  - arithmetic operator (modulo) 277
  - wildcard character 295
- . (period)
  - preceding milliseconds 65, 133
  - separator for qualifier names 290
- | (pipe) “or” bitwise operator 278
- + (plus)
  - arithmetic operator 277
  - in integer data 13
  - null values and 280
  - string concatenation operator 279
- £ (pound sterling sign)
  - in identifiers 287
  - in money datatypes 19
- “ ” (quotation marks)
  - comparison operators and 280
  - enclosing constant values 67
  - enclosing *datetime* values 21
  - enclosing empty strings 283, 285
  - in expressions 285
  - literal specification of 285
- / (slash) arithmetic operator (division) 277
- [ ] (square brackets)
  - character set wildcard 295, 296
  - in SQL statements xvii
- [^] (square brackets and caret) character set wildcard 295
- ~ (tilde) “not” bitwise operator 278
- \_ (underscore)
  - object identifier prefix 263, 286
  - in temporary table names 288
  - character string wildcard 295, 296
- ¥ (yen sign)
  - in identifiers 287
  - in money datatypes 19
- @@*cursor\_rows* global variable 268

## Numerics

“0x” prefix 30, 32  
 21st century numbers 21

## A

abbreviations

**chars** for **characters**, **patindex** 190, 194  
 date parts 64, 132

**abort** option, **lct\_admin** function 165

**abs** mathematical function 70

accent sensitivity, wildcard characters and 295

ACF. *See* Application Context Facility

**acos** mathematical function 71

adding

interval to a date 125  
 timestamp column 253  
 user-defined datatypes 42

addition operator (+) 277

aggregate functions 49–55

*See also* row aggregates; *individual function names*

**avg** 76

**count** 111

**count\_big** 113–114

difference from row aggregates 53

**group by** clause and 50, 51

**having** clause and 50

**max** 178

**min** 180

scalar aggregates 50

**sum** 240

vector aggregates 50

aggregate functions and cursors 53

**all** keyword including subqueries 281

**alter table** command, adding *timestamp* column 253

ampersand (&) “and” bitwise operator 278

and (&) bitwise operator 278

**and** keyword

in expressions 283

range-end 281

angles, mathematical functions for 71

ANSI SQL datatypes 11

**any** keyword in expressions 281

application attributes 217

Application Context Facility (ACF) 217

application contexts

getting 147

listing 172

removing 209

setting 217

approximate numeric datatypes 16

**arithabort** option, **set**

**arith\_overflow** and 11, 61

mathematical functions and **arith\_overflow** 66

mathematical functions and **numeric\_truncation**  
 62, 66

**arithignore** option, **set**

**arith\_overflow** and 61

mathematical functions and **arith\_overflow** 66

arithmetic

errors 66

expressions 276

operations, approximate numeric datatypes and 16

operations, exact numeric datatypes and 13

operations, money datatypes and 18

operators, in expressions 277

ASCII characters 72

**ascii** string function 72

**asin** mathematical function 73

asterisk (\*)

multiplication operator 277

overlength numbers 232

**atan** mathematical function 74

**@@authmech** global variable 267

**@@bootcount** global variable 267

**@@boottime** global variable 267

**@@bulkarraysize** global variable 267

**@@bulkbatchsize** global variable 267

**@@char\_convert** global variable 268

**@@cis\_rpc\_handling** global variable 268

**@@cis\_version** global variable 268

**@@client\_csexpansion** global variable 268

**@@client\_csid** global variable 268

**@@client\_csname** global variable 268

**@@cmpstate** global variable 268

**@@connections** global variable 268

**@@cpu\_busy** global variable 268

**@@curluid** global variable 268

**@@datefirst** global variable 268

**@@dbts** global variable 268

**@@error** global variable 268

- @@errorlog global variable 269
- @@failedoverconn global variable 269
- @@fetch\_status global variable 269
- @@guestuserid global variable 269
- @@hacmpservername global variable 269
- @@haconnection global variable 269
- @@heapmemsize global variable 269
- @@identity global variable 269
- @@idle global variable 269
- @@invaliduserid global variable 269
- @@io\_busy global variable 269
- @@isolation global variable 269
- @@kernel\_addr global variable 269
- @@kernel\_size global variable 269
- @@langid global variable 269
- @@language global variable 269
- @@lock\_timeout global variable 269
- @@max\_connections global variable 269
- @@max\_precision global variable 269
- @@maxcharlen global variable 269
- @@maxgroupid global variable 269
- @@maxpagesize global variable 269
- @@maxspid global variable 269
- @@maxsuid global variable 270
- @@maxuserid global variable 270
- @@mempool\_addr global variable 270
- @@min\_poolsize global variable 270
- @@mingroupid global variable 270
- @@minspid global variable 270
- @@minsuid global variable 270
- @@minuserid global variable 270
- @@monitors\_active global variable 270
- @@ncharsize global variable 270
- @@nestlevel global variable 270
- @@nodeid global variable 270
- @@optgoal global variable 270
- @@options global variable 270
- @@optimeout global variable 270
- @@pack\_received global variable 270
- @@pack\_sent global variable 270
- @@packet\_errors global variable 270
- @@pagesize global variable 270
- @@parallel\_degree global variable 270
- @@probesuid global variable 270
- @@procid global variable 270
- @@recovery\_state global variable 271
- @@repartition\_degree global variable 271
- @@resource\_granularity global variable 271
- @@rowcount global variable 271
- @@scan\_parallel\_degree global variable 271
- @@servername global variable 271
- @@setrowcount global variable 271
- @@shmem\_flags global variable 271
- @@spid global variable 271
- @@sqlstatus global variable 271
- @@ssl\_ciphersuite global variable 272
- @@stringsize global variable 272
- @@tempdbid global variable 272
- @@textcolid global variable 39, 272
- @@textdatptnid global variable 272
- @@textdbid global variable 39, 272
- @@textobjid global variable 39, 272
- @@textptnid global variable 272
- @@textptr global variable 39, 272
- @@textptr\_parameters global variable 272
- @@textsize global variable 39, 272
- @@texts global variable 39, 272
- @@thresh\_hysteresis global variable 272
- @@timeticks global variable 272
- @@total\_errors global variable 272
- @@total\_read global variable 272
- @@total\_write global variable 272
- @@tranchained global variable 272
- @@trancount global variable 272
- @@transactional\_rpc global variable 272
- @@transtate global variable 273
- @@unicharsize global variable 273
- @@version global variable 273
- @@version\_as\_integer global variable 273
- @@version\_number global variable 273
- atan2** mathematical function 75
- attributes, setting in an application 217
- audit\_event\_name** function 78
- auditing
  - audit\_event\_name** function 78
- @@authmech global variable 267
- automatic operations, updating columns with *timestamp* 19
- avg** aggregate function 76

## B

backslash (\) for character string continuation 285  
 Backus Naur Form (BNF) notation xvii  
 base 10 logarithm function 175  
**between** keyword 281  
*bigint* datatype 13  
**biginttohex** datatype conversion function 80  
 binary  
   datatypes 30–32  
   datatypes, “0x” prefix 30, 32  
   datatypes, trailing zeros in 31  
   expressions 275  
   expressions, concatenating 279  
   representation of data for bitwise operations 278  
   sort 101, 226  
*binary* datatype 30–33  
*bit* datatype 33  
 bitwise operators 278–279  
 blanks  
   *See also* spaces, character  
   character datatypes and 27–30  
   comparisons 280  
   empty string evaluated as 285  
   **like** and 296  
   removing leading, with **ltrim** function 177  
   removing trailing, with **rtrim** function 216  
 BNF notation in SQL statements xvii  
 boolean (logical) expressions 275  
 @@*bootcount* global variable 267  
 @@*boottime* global variable 267  
 brackets. *See* square brackets [ ]  
 browse mode and *timestamp* datatype 19, 252  
 built-in function, ACF 217  
 built-in functions 43–264  
   *See also individual function names*  
   aggregate 49  
   conversion 55  
   date 64  
   image 69  
   mathematical 65  
   security 66  
   string 67  
   system 68  
   text 69  
   type conversion 103–108  
 @@*bulkarraysize* global variable 267

@@*bulkbatchsize* global variable 267  
**by** row aggregate subgroup 53

## C

calculating dates 128  
**caldayofweek** date part 132  
**calweekofyear** date part 132  
**calyearofweek** date part 132  
**case** expressions 81–83, 186–187  
   null values and 82, 94, 186  
 case sensitivity  
   comparison expressions and 280, 295  
   identifiers and 288  
   in SQL xviii  
**cast** function 84–86  
**cdw**. *See caldayofweek* date part  
**ceiling** mathematical function 87  
 chains of pages, *text* or *image* data 35  
*char* datatype 25–27  
   in expressions 284  
**char** string function 89  
 @@*char\_convert* global variable 268  
**char\_length** string function 91  
 character data, avoiding “NULL” in 283  
 character datatypes 25–30  
 character expressions  
   blanks or spaces in 27–30  
   defined 275  
   syntax 276  
 character sets  
   conversion errors 293  
   iso\_1 293  
   multibyte 292  
   object identifiers and 292  
 character strings  
   continuation with backslash (\) 285  
   empty 285  
   specifying quotes within 285  
   wildcards in 293  
 characters  
   *See also* spaces, character  
   “0x” 30, 32  
   0x 62  
   deleting, using **stuff** function 237



- number of 91
- wildcard 293–299
- charindex** string function 93
- @@cis\_rpc\_handling** global variable 268
- @@cis\_version** global variable 268
- client, host computer name and 156
- @@client\_csexpansion** global variable 268
- @@client\_csid** global variable 268
- @@client\_csname** global variable 268
- @@cmpstate** global variable 268
- coalesce** function 94–95
- coalesce** keyword, **case** 94
- codes, **soundex** 228
- col\_length** system function 96
- col\_name** system function 97
- colon (:), preceding milliseconds 133
- column identifiers. *See* identifiers.
- column name
  - as qualifier 290
  - in parentheses 53
  - returning 97
- columns
  - identifying 290
  - length definition 96
  - length of 96
  - numeric, and row aggregates 53
  - sizes of (list) 2
- comma (,)
  - default print format for money values 18
  - not allowed in money values 19
  - in SQL statements xvii
- compare** system function 98
- comparing values
  - difference** string function 143
  - in expressions 280
  - timestamp* 252
- comparison operators
  - See also* relational expressions
  - in expressions 280
  - symbols for 280
- compute** clause and row aggregates 52
- computing dates 128
- concatenation
  - null values 280
  - using + operator 279
  - using || operator 279
- @@connections** global variable 268
- constants
  - and string functions 67
  - comparing in expressions 284
  - expression for 275
  - string functions and 67
- continuation lines, character string 285
- conventions
  - See also* syntax
  - identifier name 290
  - Transact-SQL syntax xvii
  - used in the Reference Manual xvi
- conversion
  - automatic values 9
  - between character sets 293
  - character value to ASCII code 72
  - dates used with **like** keyword 24
  - degrees to radians 201
  - implicit 9, 284
  - integer value to character value 89, 250
  - lower to higher datatypes 284
  - lowercase to uppercase 254, 255, 256, 257
  - null values and automatic 10
  - radians to degrees 139
  - string concatenation 279
  - styles for dates 104
  - uppercase to lowercase 176
- convert** datatype conversion function 103
  - concatenation and 279
  - date styles 104
- converting hexadecimal numbers 62
- cos** mathematical function 109
- cot** mathematical function 110
- count** aggregate function 111
- count\_big** aggregate function 113–114
- CP 850 Alternative
  - lower case first 101, 226
  - no accent 101, 226
  - no case preference 101, 226
- CP 850 Scandinavian
  - dictionary 101, 226
- @@cpu\_busy** global variable 268
- create table** command and null values 283
- @@curlroid** global variable 268
- curly braces ({}), in SQL statements xvii
- currency symbols 19, 287

## Index

current user  
  roles of 219  
  **suser\_id** system function 242  
  **suser\_name** system function 243  
  **user\_id** system function 261  
  **user\_name** system function 262  
**current\_date** date function 115  
**current\_time** date function 116  
cursors and aggregate functions 53  
**curunreservedpgs** system function 117  
**cwk**. *See* **calweekofyear** date part  
**cyr**. *See* **calyearofweek** date part  
cyrillic characters 292

## D

**data\_pages** system function 119–120  
database object owners and identifiers 291  
database objects  
  *See also individual object names*  
  ID number 188  
  identifier names 285  
  user-defined datatypes as 42  
database owners  
  name as qualifier 290, 291  
  objects and identifiers 291  
databases  
  *See also database objects*  
  getting name of 138  
  ID number, **db\_id** function 137  
**datachange** system function 121–122  
**datalength** system function 123  
  compared to **col\_length** 96  
datatype conversions  
  **biginttohex** 80  
  binary and numeric data 63  
  bit information 63  
  character information 58  
  **convert** function 103, 106  
  date and time information 60  
  domain errors 62, 85, 106  
  functions for 55–63  
  hexadecimal-like information 62  
  **hextobigint** 153  
  **hextoint** 154

**hextoint** function 153, 154  
  *image* 63, 85, 107  
  implicit 56  
  **inttohex** 160  
  money information 59  
  numeric information 59, 60  
  overflow errors 61  
  rounding during 59  
  scale errors 61  
datatype precedence. *See* precedence  
datatypes 1–42  
  *See also user-defined datatypes; individual datatype names*  
  ANSI SQL 11  
  approximate numeric 16  
  binary 30–32  
  *bit* 33  
  date and time 20–24  
  *datetime* values comparison 280  
  decimal 14–15  
  dropping user-defined 42  
  exact numeric 12–15  
  hierarchy 7  
  integer 13–14  
  mixed, arithmetic operations on 277  
  summary of 2–4  
  synonyms for 2  
  trailing zeros in *binary* 31  
  Transact-SQL extensions 11  
  user-defined 11  
  *varbinary* 224  
*date and time* datatype 21–25  
*date* datatype 20  
date functions 64–65  
  *See also individual function names*  
  **current\_date** 115  
  **current\_time** 116  
  **dateadd** 124  
  **datediff** 127  
  **datename** 130  
  **datepart** 132  
  **day** 136  
  **getdate** 149  
  **month** 181  
  **year** 265  
date parts

- abbreviation names and values 64, 132
  - caldayofweek** 132
  - calweekofyear** 132
  - calyearofweek** 132
  - entering 21
  - order of 22
  - dateadd** date function 124
  - datediff** date function 127
  - datediff** function 128
  - datefirst** option, **set** 130, 135
  - dateformat** option, **set** 22
  - datename** date function 130
  - datepart** date function 132
  - dates
    - comparing 280
    - datatypes 20–24
    - default display settings 23
    - display formats 20
    - earliest allowed 21, 64, 125
    - entry formats 22
    - pre-1753 datatypes for 64, 125
  - datetime* datatype 21–25
    - comparison of 280
    - conversion 24
    - date functions and 133
    - values and comparisons 24
  - day** date function 136
  - day** date part 64, 132
  - dayofyear** date part abbreviation and values 64, 132
  - db\_id** system function 137, 138
  - db\_name** system function 138
  - DB-Library programs, overflow errors in 77, 241
  - @@dbts** global variable 268
  - dd**. *See* **day** date part.
  - decimal* datatype 14–15
  - decimal numbers
    - round** function and 213
    - str** function, representation of 232
  - decimal points
    - datatypes, allowing in 14
    - in integer data 13
  - default settings
    - date display format 20, 23
    - weekday order 135
  - default values
    - datatype length 103
    - datatype precision 103
    - datatype scale 103
  - degrees** mathematical function 139
  - degrees, conversion to radians 201
  - delete** command and *text* row 38
  - derived\_stat** system function 140
  - devices. *See* *sysdevices* table.
  - difference** string function 143
  - division operator (*/*) 277
  - dollar sign (\$)
    - in identifiers 287
    - in money datatypes 19
  - domain rules, mathematical functions errors in 66
  - dots (..) for omitted name elements 291
  - double pipe (||)
    - string concatenation operator 279
  - double precision* datatype 17
  - double-byte characters. *See* Multibyte character sets.
  - double-precision floating-point values 17
  - doubling quotes
    - in expressions 285
    - in character strings 28
  - dropping
    - character with **stuff** function 237
    - leading or trailing blanks 177
  - duplicate rows, *text* or *image* 41
  - duplication of text. *See* **replicate** string function
  - dw**. *See* **weekday** date part.
  - dy**. *See* **dayofyear** date part.
- ## E
- e or E exponent notation
    - approximate numeric datatypes 17
    - float* datatype 6
    - money datatypes 18
  - embedded spaces. *See* spaces, character.
  - empty string (“ ”) or (‘ ’)
    - not evaluated as null 283
    - as a single space 30, 285
  - enclosing quotes in expressions 285
  - equal to. *See* comparison operators
  - @@error** global variable 268
  - error handling, domain or range 66
  - @@errorlog** global variable 269

## Index

- errors
    - arithmetic overflow 61
    - cast** function 85
    - convert** function 58–62, 106
    - divide-by-zero 61
    - domain 62, 85, 106
    - scale 61
    - trapping mathematical 66
  - escape characters 298
  - escape** keyword 298–299
  - european characters in object identifiers 293
  - exact numeric datatypes 12–15
    - arithmetic operations and 13
  - exists** keyword in expressions 281
  - exp** mathematical function 144
  - explicit null value 283
  - exponent, datatype (e or E)
    - approximate numeric types 17
    - float* datatype 6
    - money types 18
  - exponential value 144
  - expressions
    - defined 275
    - enclosing quotes in 285
    - including null values 281
    - name and table name qualifying 291
    - types of 275
- ## F
- @@failedoverconn** global variable 269
  - @@fetch\_status** global variable 269
  - finding
    - database ID 137
    - database name 138
    - server user ID 242
    - server user name 243, 244, 252, 258
    - starting position of an expression 93
    - user aliases 264
    - user IDs 261
    - user names 260, 262
    - valid identifiers 263
  - first-of-the-months, number of 128
  - fixed-length columns
    - binary datatypes for 30
    - character datatypes for 26
    - null values in 10
  - float* datatype 17
  - floating-point data 275
    - str** character representation of 232
  - floor** mathematical function 145, 146
  - formats, date. *See* dates.
  - free pages, **curunreservedpgs** system function 118
  - front-end applications, browse mode and 252
  - functions 43
    - abs** mathematical function 70
    - acos** mathematical function 71
    - aggregate 49
    - ascii** string function 72
    - asin** mathematical function 73
    - atan** mathematical function 74
    - atn2** mathematical function 75
    - avg** aggregate function 76
    - biginttohex** datatype conversion function 80
    - cast** function 84–86
    - ceiling** mathematical function 87
    - char** string function 89
    - char\_length** string function 91
    - charindex** string function 93
    - coalesce** function 94–95
    - col\_length** system function 96
    - col\_name** system function 97
    - compare** system function 98
    - conversion 55
    - convert** datatype conversion function 103
    - cos** mathematical function 109
    - cot** mathematical function 110
    - count** aggregate function 111
    - count\_big** aggregate function 113–114
    - current\_date** date function 115
    - current\_time** date function 116
    - curunreservedpgs** system function 117
    - data\_pages** system function 119–120
    - datachange** system function 121–122
    - datalength** system function 123
    - date 64
    - dateadd** date function 124
    - datediff** date function 127
    - datetime** date function 130
    - datepart** date function 132
    - day** date function 136

**db\_id** system function 137, 138  
**degrees** mathematical function 139  
**derived\_stat** system function 140  
**difference** string function 143  
**exp** mathematical function 144  
**floor** mathematical function 145  
**get\_appcontext** security function 147  
**getdate** date function 149  
**has\_role** system function 151  
**hextobigint** datatype conversion function 153  
**hextoint** datatype conversion function 154  
**host\_id** system function 155  
**host\_name** system function 156  
 image 69  
**index\_col** system function 158  
**index\_colorder** system function 159  
**inttohex** datatype conversion function 160  
**is\_quiesced** function 161–162  
**is\_sec\_service\_on** security function 163  
**isnull** system function 164  
**lct\_admin** system function 165  
**left** system function 168  
**len** string function 170  
**license\_enabled** system function 171  
**list\_appcontext** security function 172  
**lockscheme** system function 173  
**log** mathematical function 174  
**log10** mathematical function 175  
**lower** string function 176  
**ltrim** string function 177  
 mathematical 65  
**max** aggregate function 178  
**min** aggregate function 180  
**month** date function 181  
**mut\_excl\_roles** system function 182  
**newid** system function 183  
**next\_identity** system function 185  
**object\_id** system function 188  
**object\_name** system function 189  
**pagesize** system function 190  
**partition\_id** 192  
**partition\_id** system function 192  
**partition\_name** 193  
**partition\_name** system function 193  
**patindex** string function 194  
**pi** mathematical function 197  
**power** mathematical function 198  
**proc\_role** system function 199  
**radians** mathematical function 201  
**rand** mathematical function 202  
**replicate** string function 203  
**reserved\_pages** system function 204  
**reverse** string function 206  
**right** string function 207  
**rm\_appcontext** security function 209  
**role\_contain** system function 210  
**role\_id** system function 211  
**role\_name** system function 212  
**round** mathematical function 213  
**row\_count** system function 215  
**rtrim** string function 216  
 security 66  
**set\_appcontext** security function 217  
**show\_role** system function 219  
**show\_sec\_services** security function 220  
**sign** mathematical function 221  
**sin** mathematical function 222  
**sortkey** 224  
**sortkey** system function 223  
**soundex** string function 228  
**space** string function 229  
**sqrt** mathematical function 231  
**square** mathematical function 230  
**str** string function 232  
**str\_replace** string function 234  
 string 67  
**stuff** string function 236  
**substring** string function 238  
**sum** aggregate function 240  
**suser\_id** system function 242  
**suser\_name** system function 243  
**syb\_quit** system function 244  
**syb\_sendmsg** 245  
 system 68  
**tan** mathematical function 246  
**tempdb\_id** system function 247  
 text 69  
**textptr** text and image function 248  
**textvalid** text and image function 249  
**to\_unichar** string function 250  
**tran\_dumptable\_status** string function 251  
**tsequal** system function 252

**uhighsurr** string function 254  
**ulowsurr** string function 255  
**upper** string function 256  
**uscalar** string function 257  
**used\_pages** system function 258  
**user** system function 260  
**user\_id** system function 261  
**user\_name** system function 262  
**valid\_name** system function 263  
**valid\_user** system function 264  
**year** date function 265  
 functions, built-in, type conversion 103–108

## G

GB Pinyin 101, 226  
**get\_appcontext** security function 147  
**getdate** date function 149  
**getutcddate** to obtain the GMT 150  
 global variables  
   @@*authmech* 267  
   @@*bootcount* 267  
   @@*boottime* 267  
   @@*bulkarraysize* 267  
   @@*bulkbatchsize* 267  
   @@*char\_convert* 268  
   @@*cis\_rpc\_handling* 268  
   @@*cis\_version* 268  
   @@*client\_csexpansion* 268  
   @@*client\_csid* 268  
   @@*client\_csname* 268  
   @@*cmpstate* 268  
   @@*connections* 268  
   @@*cpu\_busy* 268  
   @@*curluid* 268  
   @@*cursor\_rows* 268  
   @@*dbts* 268  
   @@*error* 268  
   @@*errorlog* 269  
   @@*failedoverconn* 269  
   @@*fetch\_status* 269  
   @@*guestuserid* 269  
   @@*hacmpservername* 269  
   @@*hacconnection* 269  
   @@*heapmemsize* 269  
   @@*identity* 269  
   @@*idle* 269  
   @@*invaliduserid* 269  
   @@*io\_busy* 269  
   @@*isolation* 269  
   @@*kernel\_addr* 269  
   @@*kernel\_size* 269  
   @@*langid* 269  
   @@*language* 269  
   @@*lock\_timeout* 269  
   @@*max\_connections* 269  
   @@*max\_precision* 269  
   @@*maxcharlen* 269  
   @@*maxgroupid* 269  
   @@*maxpagesize* 269  
   @@*maxspid* 269  
   @@*maxsuid* 270  
   @@*maxuserid* 270  
   @@*mempool\_addr* 270  
   @@*min\_poolsize* 270  
   @@*mingroupid* 270  
   @@*minspid* 270  
   @@*minsuid* 270  
   @@*minuserid* 270  
   @@*monitors\_active* 270  
   @@*ncharsize* 270  
   @@*nestlevel* 270  
   @@*nodeid* 270  
   @@*optgoal* 270  
   @@*options* 270  
   @@*opttimeout* 270  
   @@*pack\_received* 270  
   @@*pack\_sent* 270  
   @@*packet\_errors* 270  
   @@*pagesize* 270  
   @@*parallel\_degree* 270  
   @@*probesuid* 270  
   @@*procid* 270  
   @@*recovery\_state* 271  
   @@*repartition\_degree* 271  
   @@*resource\_granularity* 271  
   @@*rowcount* 271  
   @@*scan\_parallel\_degree* 271  
   @@*servername* 271  
   @@*setrowcount* 271  
   @@*shmem\_flags* 271

@@*spid* 271  
 @@*sqlstatus* 271  
 @@*ssl\_ciphersuite* 272  
 @@*stringsize* 272  
 @@*tempdbid* 272  
 @@*textcolid* 272  
 @@*textdatamid* 272  
 @@*textdbid* 272  
 @@*textobjid* 272  
 @@*textptnid* 272  
 @@*textptr* 272  
 @@*textptr\_parameters* 272  
 @@*textsize* 272  
 @@*textts* 272  
 @@*thresh\_hysteresis* 272  
 @@*timeticks* 272  
 @@*total\_errors* 272  
 @@*total\_read* 272  
 @@*total\_write* 272  
 @@*tranchained* 272  
 @@*trancount* 272  
 @@*transactional\_rpc* 272  
 @@*transtate* 273  
 @@*unicharsize* 273  
 @@*version* 273  
 @@*version\_as\_integer* 273  
 @@*version\_number* 273  
 @@*datefirst* 268  
 greater than. *See* comparison operators.  
 Greek characters 292  
**group by** clause and aggregate functions 50, 51  
 guest users 261  
 @@*guestuserid* global variable 269

## H

@@*hacmpservername* global variable 269  
 @@*haconnection* global variable 269  
**has\_role** system function 151  
**having** clause and aggregate functions 50  
 @@*heapmemsize* global variable 269  
 hexadecimal numbers, converting 62  
**hextobigint** datatype conversion function 153  
**hextoint** datatype conversion function 154  
**hextoint** function 153, 154

**hh.** *See* **hour** date part.  
 hierarchy  
     *See also* precedence  
         operators 277  
 historic dates, pre-1753 64, 125  
 host computer name 156  
 host process ID, client process 155  
**host\_id** system function 155  
**host\_name** system function 156  
**hour** date part 64, 132

## I

identifiers 285–293  
     case sensitivity and 288  
     long 285  
     renaming 292  
     short 287  
     system functions and 263  
 identities  
     **sa\_role** and Database Owner 261  
     server user (**suser\_id**) 243  
     user (**user\_id**) 261  
 @@*identity* global variable 269  
**identity\_burn\_max** function 157  
 @@*idle* global variable 269  
 IDs, server role and **role\_id** 211  
 IDs, user  
     database (**db\_id**) 137  
     server user 243  
     **user\_id** function for 242  
*image* datatype 34–41  
     initializing 36  
     null values in 37  
     prohibited actions on 39  
 image functions 69  
 implicit conversion of datatypes 9, 284  
**in** keyword in expressions 281  
**index\_col** system function 158  
**index\_colorder** system function 159  
 indexes  
     *See also* clustered indexes; database objects;  
         nonclustered indexes  
     *sysindexes* table 37  
 initializing *text* or *image* columns 38

- inserting
    - automatic leading zero 32
    - spaces in text strings 229
  - int* datatype 13
    - aggregate functions and 77, 241
  - integer data in SQL 275
  - integer datatypes, converting to 62
  - integer remainder. *See* Modulo operator (%)
  - internal datatypes of null columns 10
    - See also* datatypes
  - internal structures, pages used for 204
  - inttohex** datatype conversion function 160
  - @@invaliduserid** global variable 269
  - @@io\_busy** global variable 269
  - is not null** keyword in expressions 281
  - is\_quiesced** function 161–162
  - is\_sec\_service\_on** security function 163
  - isnull** system function 164
  - ISO 8859-5 Cyrillic dictionary 101, 227
  - ISO 8859-5 Russian dictionary 101, 227
  - ISO 8859-9 Turkish dictionary 101, 227
  - iso\_1 character set 293
  - @@isolation** global variable 269
  - isql** utility command
    - See also* *Utility Guide* manual
    - approximate numeric datatypes and 17
- J**
- Japanese character sets and object identifiers 293
  - joins
    - count** or **count(\*)** with 112, 113
    - null values and 282
- K**
- @@kernel\_addr** global variable 269
  - @@kernel\_size** global variable 269
  - keywords 301–304
    - Transact-SQL 287, 301–302
- L**
- @@langid** global variable 269
  - @@language** global variable 269
  - languages, alternate
    - effect on date parts 135
    - weekday order and 135
  - last-chance threshold and **lct\_admin** function 166
  - last-chance thresholds 167
  - latin-1 English, French, German
    - dictionary 101, 226
    - no accent 101, 227
  - latin-1 Spanish
    - no accent 101, 227
    - no case 101, 227
  - lct\_admin** system function 165, 167
  - leading blanks, removal with **ltrim** function 177
  - leading zeros, automatic insertion of 32
  - left** system function 168
  - len** string function 170
  - length
    - See also* size
    - of expressions in bytes 123
    - identifiers 285
    - of columns 96
  - less than. *See* comparison operators
  - license\_enabled** system function 171
  - like** keyword
    - searching for dates with 24
    - wildcard characters used with 295
  - linkage, page. *See* pages, data
  - list\_appctx** security function 172
  - listing datatypes with types 7
  - lists
    - functions 44
  - literal character specification
    - like** match string 297
    - quotes (“ ”) 285
  - literal values
    - datatypes of 6
    - null 283
  - @@lock\_timeout** global variable 269
  - lockscheme** system function 173
  - log** mathematical function 173, 174
  - log10** mathematical function 175
  - logarithm, base 10 175
  - logical expressions 275



syntax 276  
 truth tables for 283  
**when...then** 81, 94, 186  
**log10** mathematical function 175  
*longsysname* datatype 33  
 lower and higher datatypes. *See* precedence.  
**lower** string function 176  
 lowercase letters, sort order and 288  
     *See also* case sensitivity  
**ltrim** string function 177

## M

macintosh character set 293  
 matching  
     *See also* Pattern matching  
     name and table name 291  
 mathematical functions 65  
   **abs** 70  
   **acos** 71  
   **asin** 73  
   **atan** 74  
   **atn2** 75  
   **ceiling** 87  
   **cos** 109  
   **cot** 110  
   **degrees** 139  
   **exp** 144  
   **floor** 145  
   **log** 174  
   **log10** 175  
   **pi** 197  
   **power** 198  
   **radians** 201  
   **rand** 202  
   **round** 213  
   **sign** 221  
   **sin** 222  
   **sqrt** 231  
   **square** 230  
   **tan** 246  
**max** aggregate function 178  
 @@*max\_connections* global variable 269  
 @@*max\_precision* global variable 269  
 @@*maxcharlen* global variable 269

@@*maxgroupid* global variable 269  
 @@*maxpagesize* global variable 269  
 @@*maxspid* global variable 269  
 @@*maxsuid* global variable 270  
 @@*maxuserid* global variable 270  
 @@*mempool\_addr* global variable 270  
 messages and mathematical functions 66  
**mi.** *See* **minute** date part  
 midnights, number of 128  
**millisecond** date part 65, 132  
 millisecond values, **datediff** results in 128  
**min** aggregate function 180  
 @@*min\_poolsize* global variable 270  
 @@*mingroupid* global variable 270  
 @@*minspid* global variable 270  
 @@*minsuid* global variable 270  
 minus sign (-)  
   in integer data 13  
   subtraction operator 277  
 @@*minuserid* global variable 270  
**minute** date part 65, 132  
 mixed datatypes, arithmetic operations on 277  
**mm.** *See* **month** date part  
**mm.** *See* **month** date part.  
*model* database, user-defined datatypes in 42  
 modulo operator (%) 277  
 money  
   default comma placement 18  
   symbols 287  
*money* datatype 18  
   arithmetic operations and 18  
 @@*monitors\_active* global variable 270  
**month** date function 181  
**month** date part 64, 132  
 month values and date part abbreviation 64, 132  
**ms.** *See* **millisecond** date part  
 multibyte character sets  
   converting 58  
   identifier names 292  
   *nchar* datatype for 25  
   wildcard characters and 297  
 multiplication operator (\*) 277  
**mut\_excl\_roles** system function 182  
 mutual exclusivity of roles and **mut\_excl\_roles** 182

## N

- “N/A”, using “NULL” or 283
- names
  - See also* identifiers
  - checking with **valid\_name** 292
  - date parts 64, 132
  - db\_name** function 138
  - finding similar-sounding 228
  - host computer 156
  - index\_col** and index 158
  - object\_name** function 189
  - omitted elements of (..) 291
  - qualifying database objects 290, 292
  - user\_name** function 243
  - user\_name** function 262
  - weekday numbers and 135
- naming
  - conventions 285–293
  - database objects 285–293
  - identifiers 285–293
  - user-defined datatypes 42
- national character. *See* **nchar** datatype
- natural logarithm 173, 174
- nchar** datatype 26–27
- @@ncharsize** global variable 270
- negative sign (-) in money values 19
- nesting
  - aggregate functions 51
  - string functions 67
- @@nestlevel** global variable 270
- newid** system function 183
- next\_identity** system function 185
- @@nodeid** global variable 270
- “none”, using “NULL” or 283
- not** keyword in expressions 281
- not like** keyword 294
- not null values
  - spaces in 29
- not null values in spaces 29
- null** keyword in expressions 281
- null string in character columns 237, 283
- null values
  - column datatype conversion for 29
  - default parameters as 282
  - in expressions 282
  - text* and *image* columns 37
  - null values in a **where** clause 282
  - nullif** expressions 186–187
  - nullif** keyword 186
  - number (quantity of)
    - first-of-the-months 128
    - midnights 128
    - rows in **count(\*)** 111, 113
    - Sundays 128
  - number of characters and date interpretation 24
  - numbers
    - asterisks (\*\*) for overlength 232
    - converting strings of 30
    - database ID 137
    - object ID 188
    - odd or even binary 32
    - random float 202
    - weekday names and 135
  - numeric data and row aggregates 53
  - numeric** datatype 14
  - numeric expressions 275
    - round** function for 213
  - nvarchar** datatype 27
    - spaces in 27

## O

- object names, database
  - See also* identifiers
  - user-defined datatype names as 42
- object\_id** system function 188
- object\_name** system function 189
- objects. *See* database objects; databases
- operators
  - arithmetic 277
  - bitwise 278–279
  - comparison 280
  - precedence 277
- @@optgoal** global variable 270
- @@options** global variable 270
- @@opttimeout** global variable 270
- or** keyword in expressions 283
- order
  - See also* indexes; precedence; sort order
  - of execution of operators in expressions 277
  - of date parts 22

reversing character expression 206  
 weekday numeric 135

**order by** clause 224

other users, qualifying objects owned by 292

overflow errors in DB-Library 77, 241

ownership of objects being referenced 292

## P

**@@pack\_received** global variable 270

**@@pack\_sent** global variable 270

**@@packet\_errors** global variable 270

padding, data  
   blanks and 26  
   underscores in temporary table names 288  
   with zeros 31

pages, data  
   chain of 35  
   used for internal structures 204

**@@pagesize** global variable 270

**pagesize** system function 190

**@@parallel\_degree** global variable 270

parentheses ()  
   *See also Symbols section of this index*  
   in an expression 284  
   in SQL statements xvii

**partition\_id** function 192

**partition\_name** function 193

**patindex** string function 194  
   **text/image** function 41

pattern matching 293  
   *See also String functions; wildcard characters*  
   **charindex** string function 93  
   **difference** string function 143  
   **patindex** string function 195

percent sign (%)  
   modulo operator 277  
   wildcard character 295

period (.)  
   preceding milliseconds 133  
   separator for qualifier names 290

**pi** mathematical function 197

platform-independent conversion  
   hexadecimal strings to integer values 153, 154  
   integer values to hexadecimal strings 160

plus (+)  
   arithmetic operator 277  
   in integer data 13  
   null values and 280  
   string concatenation operator 279

pointers  
   null for uninitialized *text* or *image* column 248  
   *text* and *image* page 248  
   *text* or *image* column 36

pound sterling sign (£)  
   in identifiers 287  
   in money datatypes 19

**power** mathematical function 198

precedence  
   of lower and higher datatypes 284  
   of operators in expressions 277

preceding blanks. *See* blanks; spaces, character

precision, datatype  
   approximate numeric types 17  
   exact numeric types 14  
   money types 18

**@@probesuid** global variable 270

**proc\_role** system function 199

**@@procid** global variable 270

punctuation, characters allowed in identifiers 287

## Q

**qq.** *See* **quarter** date part

qualifier names 290, 292

**quarter** date part 64, 132

quotation marks (“ ”)  
   comparison operators and 280  
   for empty strings 283, 285  
   enclosing constant values 67  
   in expressions 285  
   literal specification of 285

## R

**radians** mathematical function 201

radians, conversion to degrees 139

**rand** mathematical function 202

range

## Index

*See also* numbers; size  
of date part values 64, 132  
**datediff** results 128  
errors in mathematical functions 66  
money values allowed 18  
of recognized dates 21  
wildcard character specification of 296, 297  
range queries  
  **and** end keyword 281  
  **between** start keyword 281  
**readtext** command and *text* data initialization requirement 38  
*real* datatype 17  
**@@recovery\_state** global variable 271  
reference information  
  datatypes 1  
  reserved words 301  
  Transact-SQL functions 43  
relational expressions 276  
  *See also* comparison operators  
removing application contexts 209  
**@@repartition\_degree** global variable 271  
**replicate** string function 203  
**reserve** option, **lct\_admin** function 165  
reserved words 301–304  
  *See also* keywords  
  database object identifiers and 285, 287  
  SQL92 302  
  Transact-SQL 301–302  
**reserved\_pages** system function 204  
**@@resource\_granularity** global variable 271  
results of row aggregate operations 53  
retrieving similar-sounding words or names 228  
**reverse** string function 206  
**right** string function 207, 208  
right-justification of **str** function 233  
**rm\_appcontext** security function 209  
role hierarchies and **role\_contain** 210  
**role\_contain** system function 210  
**role\_id** system function 211  
**role\_name** system function 212  
roles  
  checking with **has\_role** 151  
  checking with **proc\_role** 199  
  showing system with **show\_role** 219  
roles, user-defined and mutual exclusivity 182

**round** mathematical function 213  
rounding 213  
  approximate numeric datatypes 17  
  *datetime* values 20, 60  
  money values 18, 59  
  **str** string function and 232  
row aggregates 53  
  **compute** and 52  
  difference from aggregate functions 53  
**row\_count** system function 215  
**@@rowcount** global variable 271  
rows, table  
  detail and summary results 53  
  row aggregates and 53  
**rtrim** string function 216  
rules. *See* database objects.

## S

scalar aggregates and nesting vector aggregates within 51  
scale, datatype 14  
  *decimal* 9  
  IDENTITY columns 14  
  loss during datatype conversion 11  
  *numeric* 9  
**@@scan\_parallel\_degree** global variable 271  
scrollable cursor  
  **@@rowcount** 268  
search conditions and *datetime* data 24  
**second** date part 65, 132  
seconds, **datediff** results in 128  
security functions 66  
  **get\_appcontext** 147  
  **is\_sec\_service\_on** 163  
  **list\_appcontext** 172  
  **rm\_appcontext** 209  
  **set\_appcontext** 217  
  **show\_sec\_services** 220  
seed values and **rand** function 202  
**select** command 224  
  aggregates and 50  
  **for browse** 252  
  restrictions in standard SQL 51  
  in Transact-SQL compared to standard SQL 51

- select into** command not allowed with **compute** 55
- server user name and ID
  - suser\_id** function 242
  - suser\_name** function for 243
- @@servername** global variable 271
- set\_appcontext** security function 217
- @@setrowcount** global variable 271
- setting application context 217
- shift-JIS binary order 102, 227
- @@shmem\_flags** global variable 271
- short identifiers 287
- show\_role** system function 219
- show\_sec\_services** security function 220
- sign** mathematical function 221
- similar-sounding words. *See* **soundex** string function
- sin** mathematical function 222
- single quotes. *See* quotation marks
- single-byte character sets, *char* datatype for 25
- size
  - See also* length; number (quantity of); range; size limit; space allocation
  - column 96
  - floor** mathematical function 146
  - identifiers (length) 286
  - image* datatype 34
  - of **pi** 197
  - text* datatype 34
- size limit
  - approximate numeric datatypes 17
  - binary* datatype 31
  - char* columns 26
  - datatypes 2
  - double precision* datatype 17
  - exact numeric datatypes 13
  - fixed-length columns 26
  - float* datatype 17
  - image* datatype 31
  - integer value smallest or largest 146
  - money datatypes 18
  - nchar* columns 27
  - nvarchar* columns 27
  - real* datatype 17
  - varbinary* datatype 31
  - varchar* columns 26
- slash (/) division operator 277
- smalldatetime* datatype 21
  - date functions and 133
- smallint* datatype 13
- smallmoney* datatype 18
- sort order
  - character collation behavior 223, 224
  - comparison operators and 280
- sortkey** function 224
- sortkey** system function 223
- soundex** string function 228
- sp\_bindefault** system procedure and user-defined datatypes 42
- sp\_bindrule** system procedure and user-defined datatypes 42
- sp\_help** system procedure 42
- space** string function 229
- spaces, character
  - See also* blanks
  - in character datatypes 27–30
  - empty strings (“ ”) or (‘ ’) as 283, 285
  - inserted in text strings 229
  - like** *datetime* values and 24
  - not allowed in identifiers 287
- speed (Server)
  - binary* and *varbinary* datatype access 31
- @@spid** global variable 271
- SQL (used with Sybase databases). *See* Transact-SQL
- SQL standards
  - aggregate functions and 51
  - concatenation and 280
- SQLSTATE codes 305–311
  - exceptions 306–311
- @@sqlstatus** global variable 271
- sqrt** mathematical function 231
- square brackets [ ]
  - caret wildcard character [^] and 295, 297
  - in SQL statements xvii
  - wildcard specifier 295
- square** mathematical function 230
- square root mathematical function 231
- ss**. *See* **second** date part
- @@ssl\_ciphersuite** global variable 272
- storage management for *text* and *image* data 37
- str** string function 232
- str\_replace** string function 234
- string functions 67
  - See also* *text* datatype

## Index

- ascii** 72
- char** 89
- char\_length** 91
- charindex** 93
- difference** 143
- len** 170
- lower** 176
- ltrim** 177
- patindex** 194
- replicate** 203
- reverse** 206
- right** 207
- rtrim** 216
- soundex** 228
- space** 229
- str** 232
- str\_replace** 234
- stuff** 236
- substring** 238
- to\_unichar** 250
- tran\_dumpstable\_status** 251
- uhighsurr** 254
- ulowsurr** 255
- upper** 256
- uscalar** 257
- strings, concatenating 279
- @@*stringsize* global variable 272
- stuff** string function 236, 237
- style values, date representation 104
- subqueries
  - any** keyword and 281
  - in expressions 281
- substring** string function 238
- subtraction operator (-) 277
- sum** aggregate function 240
- sundays, number value 128
- suser\_id** system function 242
- suser\_name** system function 243
- syb\_quit** system function 244
- syb\_sendmsg** function 245
- symbols
  - See also* wildcard characters; *Symbols section of this index*
  - arithmetic operator 277
  - comparison operator 280
  - in identifier names 287
  - matching character strings 295
  - money 287
  - in SQL statements xvii
  - wildcards 295
- synonyms and **chars** and **characters**, **patindex** 194
- synonyms for datatypes 2
- synonyms, **chars** and **characters**, **patindex** 190
- syntax conventions, Transact-SQL xvii
- syscolumns* table 33
- sysindexes* table and *name* column in 37
- sysname* datatype 33
- sys srvroles* table and **role\_id** system function 211
- system datatypes. *See* datatypes
- system functions 68
  - col\_length** 96
  - col\_name** 97
  - compare** 98
  - curunreservedpgs** 117
  - data\_pages** 119–120
  - datachange** 121–122
  - datalength** 123
  - db\_id** 137, 138
  - derived\_stat** 140
  - has\_role** system function 151
  - host\_id** 155
  - host\_name** 156
  - index\_col** 158
  - index\_colorder** 159
  - isnull** 164
  - lct\_admin** 165
  - left** 168
  - license\_enabled** 171
  - lockscheme** 173
  - mut\_excl\_roles** 182
  - newid** system function 183
  - next\_identity** 185
  - object\_id** 188
  - object\_name** 189
  - pagesize** 190
  - proc\_role** system function 199
  - reserved\_pages** 204
  - role\_contain** 210
  - role\_id** 211
  - role\_name** 212
  - row\_count** 215
  - show\_role** 219

- sortkey** 223
  - suser\_id** 242
  - suser\_name** 243
  - syb\_quit** 244
  - tempdb\_id** 247
  - tsequal** 252
  - used\_pages** 258
  - user** 260
  - user\_id** 261
  - user\_name** 262
  - valid\_name** 263
  - valid\_user** 264
  - system roles and **show\_role** and 219
  - system tables and *sysname* datatype 33
- T**
- table pages
    - See also* pages, data
  - tables
    - identifying 290
    - names as qualifiers 290
    - worktables 50
  - tan** mathematical function 246
  - tangents, mathematical functions for 246
  - tempdb* database, user-defined datatypes in 42
  - @@tempdbid** global variable 272
  - tempdb\_id** system function 247
  - tempdbs and **tempdb\_id** system function 247
  - temporary tables, naming 288
    - number of bytes 288
    - padding 288
    - sysobjects 288
  - text and image functions
    - textptr** 248
    - textvalid** 249
  - text* datatype 34–41
    - convert** command 40
    - converting 59
    - initializing with null values 36
    - null values 37
    - prohibited actions on 39
  - text datatype and **ascii** string function 72
  - text functions 69
  - text page pointer 96
  - text pointer values 248
    - @@textcolid** global variable 39, 272
    - @@textdatptnid** global variable 272
    - @@textdbid** global variable 39, 272
    - @@textobjid** global variable 39, 272
    - @@textptnid** global variable 272
  - textptr** function 248
    - @@textptr** global variable 39, 272
  - textptr** text and image function 248
    - @@textptr\_parameters** global variable 272
    - @@textsize** global variable 39, 272
    - @@textsts** global variable 39, 272
  - textvalid** text and image function 249
  - Thai dictionary 101, 226
  - then** keyword. *See* **when...then** conditions
  - @@thresh\_hysteresis** global variable 272
  - thresholds, last-chance 167
  - time values
    - datatypes 20–24
  - timestamp* datatype 19
    - automatic update of 19
    - browse mode and 19, 252
    - comparison using **tsequal** function 252
  - @@timeticks** global variable 272
  - tinyint* datatype 13
  - to\_unichar** string function 250
    - @@total\_errors** global variable 272
    - @@total\_read** global variable 272
    - @@total\_write** global variable 272
  - trailing blanks. *See* blanks
  - tran\_dumptable\_status** string function 251
    - @@tranchained** global variable 272
    - @@trancount** global variable 272
    - @@transactional\_rpc** global variable 272
  - Transact-SQL
    - aggregate functions in 51
    - reserved words 301–302
  - Transact-SQL extensions 11
  - translation of integer arguments into binary numbers 278
    - @@transtate** global variable 273
  - triggers *See* database objects; stored procedures.
  - trigonometric functions 65, 65–246
  - true/false data, *bit* columns for 33
  - truncation
    - arithabort numeric\_truncation** 10

## Index

- binary datatypes 30
- character string 26
- datediff** results 128
- str** conversion and 233
  - temporary table names 288
- truth tables for logical expressions 283
- tsequal** system function 252
- twenty-first century numbers 21

## U

- UDP messaging 245
- uhighsurr** string function 254
- ulowsurr** string function 255
- underscore ()
  - character string wildcard 295, 296
  - object identifier prefix 263, 286
  - in temporary table names 288
- @@unicharsize** global variable 273
- unique names as identifiers 289
- unitext** datatype 34–41
- unsigned bigint** datatype 13
- unsigned int** datatype 13
- unsigned smallint** datatype 13
- updating
  - See also* changing 19
  - in browse mode 252
  - prevention during browse mode 252
- upper** string function 256, 257
- uppercase letter preference 288
  - See also* case sensitivity; **order by** clause
- us\_english language, weekdays setting 135
- uscalar** string function 257
- used\_pages** system function 258
- User Datagram Protocol messaging 245
- user IDs
  - user\_id** function for 261
  - valid\_user** function 264
- user names 262
- user names, finding 243, 262
- user objects. *See* database objects
- user** system function 260
- user\_id** system function 261
- user\_name** system function 262
- user-created objects. *See* database objects

- user-defined datatypes 11
  - See also* datatypes
  - creating 42
  - dropping 42
  - longsysname* as 33
  - sysname* as 33
- user-defined roles and mutual exclusivity 182
- using bytes** option, **patindex** string function 190, 194, 195

## V

- valid\_name** system function 263
  - using after changing character sets 292
- valid\_user** system function 264
- varbinary** datatype 30–32, 224
- varchar** datatype 27
  - datetime* values conversion to 24
  - in expressions 284
  - spaces in 27
- variable-length character. *See* **varchar** datatype
- vector aggregates 50
  - nesting inside scalar aggregates 51
- @@version** global variable 273
- @@version\_number** global variable 273
- @@version\_as\_integer** global variable 273
- view name in qualified object name 290

## W

- week** date part 64, 132
- weekday** date part 64, 132
- weekday date value, names and numbers 135
- when** keyword. *See* **when...then** conditions
- when...then** conditions 81
- where** clause, null values in a 282
- wildcard characters 293–299
  - See also* **patindex** string function
  - in a **like** match string 295
  - literal characters and 297
  - used as literal characters 297
- wk**. *See* **week** date part
- words, finding similar-sounding 228
- worktables, number of 50



**writetext** command and *text* data initialization  
requirement 38

## Y

**year** date function 265

**year** date part 64, 132

yen sign (¥)

in identifiers 287

in money datatypes 19

yes/no data, *bit* columns for 33

**yy**. *See* **year** date part

## Z

zero x (0x) 30, 32, 62

zeros, trailing, in binary datatypes 31–32

