



Using Adaptive Server Distributed Transaction
Management Features

Adaptive Server® Enterprise

15.0

DOCUMENT ID: DC31650-01-1500-01

LAST REVISED: September 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, ECRTIP, 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, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILS, smart.partners, smart.parts, smart.script, 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. 02/05

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

About This Book	v	
CHAPTER 1	Overview	1
	Distributed Transaction Management features	1
	Affected transaction types	2
	Distributed transactions coordinated by external transaction managers.....	3
	RPC and CIS transactions	4
	SYB2PC transactions.....	5
CHAPTER 2	Enabling DTM Features	7
	Installing a license key	7
	Enabling DTM features	7
	enable dtm parameter	8
	enable xact coordination parameter	8
	Configuring transaction resources.....	8
	Calculating required transaction descriptors	9
	Setting the number of transaction descriptors.....	11
CHAPTER 3	Using Adaptive Server Transaction Coordination Services	13
	Overview of transaction coordination services	13
	Hierarchical transaction coordination	14
	X/Open XA-compliant behavior in DTP environments	15
	Requirements and behavior	15
	Configuring participant server resources.....	16
	number of dtx participants parameter.....	17
	Optimizing number of dtx participants for your system	17
	Using transaction coordination services in heterogeneous environments	18
	strict dtm enforcement parameter	18
	Monitoring coordinated transactions and participants	19

CHAPTER 4

DTM Administration and Troubleshooting	21
Transactions and threads of control	21
Implications for system administrators.....	22
Lock manager changes to support detached transactions	22
Getting information about distributed transactions	23
Transaction identification in systransactions	23
Viewing active transactions with sp_transactions.....	24
Determining the commit node and gtrid with sp_transactions .	27
Crash recovery procedures for distributed transactions.....	29
Transactions coordinated with MSDTC	29
Transactions coordinated by Adaptive Server or X/Open XA..	29
Transactions coordinated with SYB2PC.....	30
Heuristically completing transactions	30
Completing prepared transactions.....	31
Completing transactions that are not prepared.....	33
Determining the commit status for Adaptive Server transactions .	
33	
Configuration considerations	35
Behavior of DDLs within distributed transactions	35

About This Book

Audience	This manual is intended for administrators or application developers who have purchased the Adaptive Server Distributed Transaction Management (DTM) feature.
How to use this book	Read this manual after you have installed Adaptive Server and its associated feature licenses.
Related documents	<p>The Sybase® Adaptive Server® Enterprise documentation set consists of the following:</p> <ul style="list-style-type: none">• 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 <i>Installation Guide</i> for your platform – describes installation, upgrade, and configuration procedures for all Adaptive Server and related Sybase products.• <i>What's New in Adaptive Server Enterprise?</i> – 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.• <i>ASE Replicator User's Guide</i> – 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.• <i>Component Integration Services User's Guide</i> – explains how to use the Adaptive Server Component Integration Services feature to connect remote Sybase and non-Sybase databases.• The <i>Configuration Guide</i> 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[®] and Adaptive Server.
 - *Java in Adaptive Server Enterprise* – describes how to install and use Java classes as data types, 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.

- *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. Available only in print version.
- *Transact-SQL User's Guide* – documents Transact-SQL, Sybase's 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's 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.	select sp_configure
Database names and datatypes are in sans serif font.	master database
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.	select <i>column_name</i> from <i>table_name</i> where <i>search_conditions</i>
Type parentheses as part of the command.	compute <i>row_aggregate</i> (<i>column_name</i>)
Double colon, equals sign indicates that the syntax is written in BNF notation. Do not type this symbol. Indicates “is defined as”.	::=
Curly braces mean that you must choose at least one of the enclosed options. Do not type the braces.	{cash, check, credit}
Brackets mean that to choose one or more of the enclosed options is optional. Do not type the brackets.	[cash check credit]
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.	cash, check, credit
The pipe or vertical bar () means you may select only one of the options shown.	cash check credit

Element	Example
An ellipsis (...) means that you can <i>repeat</i> the last unit as many times as you like.	<pre>buy thing = price [cash check credit] [, thing = price [cash check credit]]...</pre> <p>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.</p>

- 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*.

Accessibility features

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

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

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

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

If you need help

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

Overview

Topic	Page
Distributed Transaction Management features	1
Affected transaction types	2

Adaptive Server version includes several distributed transaction management features to:

- Bring Adaptive Server into full compliance with the X/Open XA protocol when acting as a resource manager, without requiring additional services such as XA-Server
- Provide support for distributed transactions coordinated by Microsoft Distributed Transaction Coordinator (MSDTC)
- Ensure consistent commit or rollback for transactions that update Adaptive Server data via remote procedure calls (RPCs) and Component Integration Services (CIS)
- Provide the framework to support additional distributed transaction management protocols in the future

This chapter presents an overview of new distributed transaction management features, and describes changes to Adaptive Server that support those features.

Distributed Transaction Management features

Adaptive Server includes these distributed transaction management features:

- Improved transaction and thread management. Adaptive Server manages all transactions as server resources, and provides the ability to attach and detach threads from transactions. These new capabilities provide a common interface for supporting clients of local server transactions, as well as clients in X/Open XA and MSDTC environments. See “Configuring transaction resources” on page 8.
- New distributed transaction coordination services. Adaptive Server provides consistent rollback and commit capabilities for transactions that modify data in remote Adaptive Servers via RPCs and CIS. New transaction coordination services guarantee the integrity of such distributed transactions, even when no external transaction manager is present. See Chapter 3, “Using Adaptive Server Transaction Coordination Services”.
- Improved recovery for prepared transactions. During recovery, Adaptive Server identifies prepared transactions that were coordinated by the X/Open XA protocol and Adaptive Server native transaction coordination services. Adaptive Server restores these transactions to the condition they were in prior to recovery, and brings the associated database online more quickly than in previous server versions. See “Crash recovery procedures for distributed transactions” on page 29.
- New dbcc commands for heuristically completing distributed transactions. See “Heuristically completing transactions” on page 30.

Affected transaction types

The new Adaptive Server DTM features affect:

- Distributed transactions coordinated by external transaction managers
- Transactions that update data using RPCs and CIS

Distributed transactions coordinated by external transaction managers

Distributed transactions can take place in an environment where an external transaction manager coordinates transaction execution using a specific protocol, such as X/Open XA. Adaptive Server supports transactions using the CICS, Encina, TUXEDO, and MSDTC transaction managers through the DTM XA interface to Adaptive Server.

Note Adaptive Server with the DTM XA interface provides features that were previously part of the XA-Server product. The XA-Server product is not required and is not included with Adaptive Server. See the *XA Interface Integration Guide for CICS, Encina, and TUXEDO* for information about the DTM XA interface.

Behavior for transaction manager-coordinated transactions

Adaptive Server natively implements several features that were part of the XA-Library and XA-Server products, and provides new recovery procedures for prepared transactions coordinated via the X/Open XA protocol. See “Configuring transaction resources” on page 8 and “Crash recovery procedures for distributed transactions” on page 29 for more information.

The XA interface to Adaptive Server has been modified to accommodate the server’s new distributed transaction management features. Changes to the XA interface are transparent to X/Open XA client applications. However, you must link Adaptive Server DTM XA interface to your X/Open XA transaction manager in order to use Adaptive Server as a resource manager. Details on all XA interface changes are described in the *XA Interface Integration Guide for CICS, Encina, and TUXEDO*.

Adaptive Server also includes support for distributed transactions coordinated by MSDTC. MSDTC clients can communicate directly with Adaptive Server using the native interface. Clients can also communicate with one or more Adaptive Server running on UNIX by using the DTM XA interface.

Note MSDTC clients using the DTM XA interface must possess `dtm_tm_role` in the Adaptive Server(s) they access. See the *XA Interface Integration Guide for CICS, Encina, and TUXEDO* for more information about `dtm_tm_role`.

RPC and CIS transactions

Local Adaptive Server transactions can update data in remote servers by using Transact-SQL remote procedure calls (RPCs) and Component Integration Services (CIS). RPC updates are accomplished by executing an RPC from within a locally-created transaction. For example:

```
sp_addserver westcoastsrv, ASEnterprise, hqsales
begin transaction rpc_tran1
update sales set commission=300 where salesid="12OZ"
exec westcoastsrv.salesdb..recordsalesproc
commit rpc_tran1
```

The above transaction updates the `sales` table on the local Adaptive Server, but also updates data on a remote server using the RPC, `recordsalesproc`.

CIS provides a way to update data on remote tables as if those tables were local. By using `sp_addobjectdef` users can create local objects in Adaptive Server that reference remote data. Updating the local object modifies data in the remote Adaptive Server. For example:

```
sp_addobjectdef salesrec,
"westcoastsrv.salesdb..sales", "table"
begin transaction cis_tran1
update sales set commission=300 where salesid="12OZ"
update salesrec set commission=300 where salesid="12OZ"
commit cis_tran1
```

New behavior for RPC and CIS transactions

Prior to Adaptive Server version 12.0, transactions that updated data via RPCs and CIS could not roll back the work of the remote server, nor could those transactions be assured that the remote work actually committed. Adaptive Server provides new transaction coordination services to assure that RPCs and CIS updates commit or roll back their work with the initiating transaction. See Chapter 3, “Using Adaptive Server Transaction Coordination Services” for more details.

If you have applications that rely on the earlier behavior of RPCs and CIS updates, you can disable transaction coordination services. See “enable xact coordination parameter” on page 8 for information.

SYB2PC transactions

SYB2PC transactions use the Sybase two-phase commit protocol to ensure that the work of a distributed transaction is committed or rolled back as a logical unit.

Adaptive Server does not modify the behavior of SYB2PC transactions. However, application developers who implement SYB2PC transactions may want to consider using Adaptive Server transaction coordination services instead. Compared to SYB2PC transactions, transactions coordinated directly by Adaptive Server use fewer network connections and execute more quickly, while still ensuring the integrity of the distributed transaction. Application code can also be simpler when Adaptive Server, rather than the application, coordinates remote transactions. See Chapter 3, “Using Adaptive Server Transaction Coordination Services” for more information.

Enabling DTM Features

This chapter describes how to enable Adaptive Server DTM Features:

Topic	Page
Installing a license key	7
Enabling DTM features	7
Configuring transaction resources	8

Installing a license key

Distributed Transaction Management is available as a separately-licensed Adaptive Server feature. Before you can enable and use DTM features, you must purchase and install a valid license for both Adaptive Server and the DTM feature.

See your *Installation Guide* for information about installing license keys and using Sybase Software Asset Management (SySAM). Contact your Sybase sales representative if you want to purchase a license for DTM or other licensed Adaptive Server features.

Enabling DTM features

After you have purchased and installed a valid license for Adaptive Server and the DTM feature, you can enable DTM features by using `sp_configure` with the `enable dtm` and `enable xact coordination` configuration parameters.

enable dtm parameter

The `enable dtm` parameter enables or disables basic DTM features. When `enable dtm` is set to 1 (on), Adaptive Server supports external transactions from MSDTC, and from X/Open XA transaction managers via the DTM XA Interface. See the *XA Interface Integration Guide for CICS, Encina, and TUXEDO* for more information.

To enable basic DTM Features, use the command:

```
sp_configure 'enable dtm', 1
```

You must restart Adaptive Server for this change to take effect.

enable xact coordination parameter

`enable xact coordination` enables or disables Adaptive Server transaction coordination services. When this parameter is enabled, Adaptive Server ensures that updates to remote Adaptive Server data commit or roll back with the original transaction. See Chapter 3, “Using Adaptive Server Transaction Coordination Services” for more information.

To enable transaction coordination, use the command:

```
sp_configure 'enable xact coordination', 1
```

You must restart Adaptive Server for this change to take effect.

Configuring transaction resources

Adaptive Server provides a common interface to support both local server transactions and external transactions coordinated by distributed transaction protocols. Distributed transaction protocol support is provided for X/Open XA, MSDTC, and native Adaptive Server transaction coordination services.

Adaptive Server manages all transactions as configurable server resources, and the System Administrator can configure the total number of resources available in a given server. Client tasks that access Adaptive Server in an X/Open XA environment can also suspend and join threads to transaction resources as needed.

This section describes how to determine and configure the total number of transaction resources available to Adaptive Server.

Calculating required transaction descriptors

Adaptive Server uses the **transaction descriptor** resource to manage transactions within a server. A transaction descriptor is an internal memory structure that Adaptive Server uses to represent a transaction.

Upon starting, Adaptive Server allocates a fixed number of transaction descriptors based on the value of the configuration parameter `txn to pss ratio` and places them in a pool. Adaptive Server obtains transaction descriptors from the pool as they are needed for new transactions. As transactions complete, descriptors are returned to the pool. If there are no transaction descriptors available, transactions may be delayed as Adaptive Server waits for descriptors to become freed.

To properly configure the number of transaction descriptors, it is important that you understand exactly when Adaptive Server tries to obtain new descriptors from the global pool. A new transaction descriptor is required when:

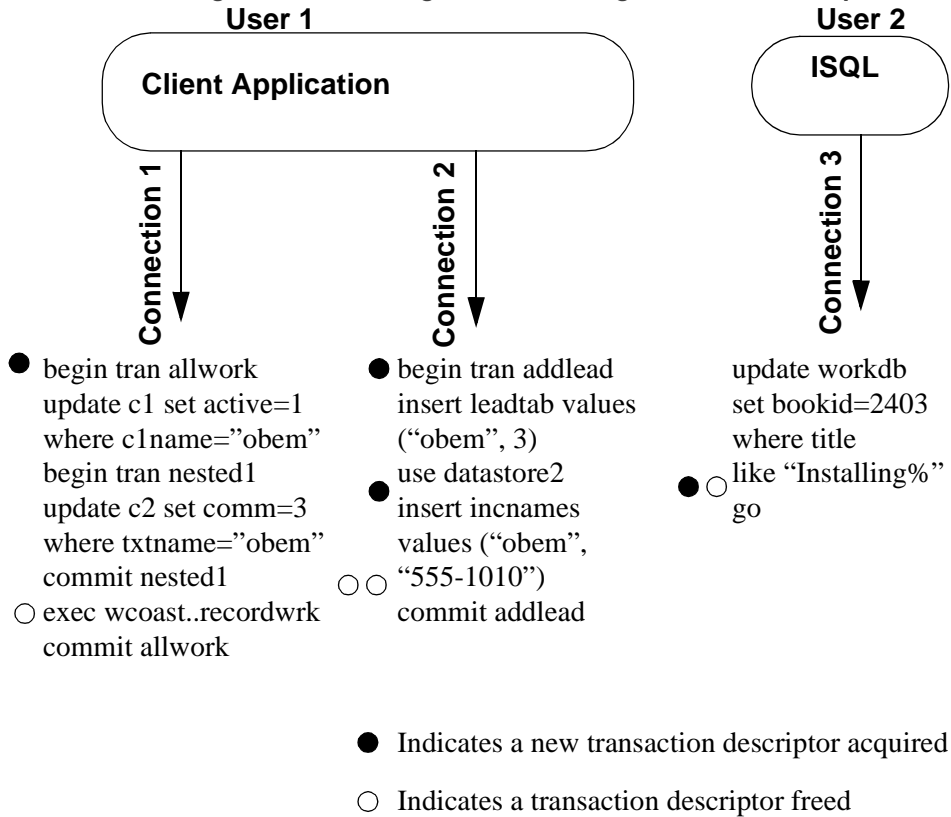
- A client connection initiates a new, outer-level transaction. This can occur explicitly, when the client executes an outer-level `begin transaction` command. It can also occur implicitly, when a client modifies data without entering a `begin transaction` command.

Once an outer-level transaction has begun, future nested `begin transaction` commands do not require additional transaction descriptors. Allocation and deallocation of the transaction descriptor is dictated by the outer-most block of the transaction.

- An existing transaction modifies a second database (a multi-database transaction). A multi-database transaction requires a dedicated transaction descriptor for *each* database it accesses.

Figure 2-1 illustrates how Adaptive Server obtains and releases transaction descriptors for different transaction types.

Figure 2-1: Allocating and deallocating transaction descriptors



In Figure 2-1, Adaptive Server uses a total of three transaction descriptors for User 1, who accesses the server through a pair of client connections. The server allocates a single descriptor for transaction `allwork`, which is freed when that transaction commits. The nested transaction, `nested1`, does not require a dedicated transaction descriptor.

Transaction `addlead`, a multi-database transaction, requires two transaction descriptors—one for the outer transaction block, and one for modifying a second database, `datastore2`. Both transaction descriptors are released when the outer transaction block commits.

User 2, accessing Adaptive Server from isql, also requires a dedicated transaction descriptor. Even though User 2 did not explicitly create an outer transaction block with `begin transaction`, Adaptive Server implicitly creates a transaction block to execute the `update` command. The transaction descriptor associated with this block is acquired after the `go` command, and released after the `insert` has completed.

Because transaction descriptors consume memory that can be used by other Adaptive Server services, it is important that you use only enough descriptors to satisfy the maximum number of transactions that may be required at any given time.

Setting the number of transaction descriptors

Once you have determined the number of transaction descriptors to use in your system, use `sp_configure` to set the value of `txn to pss ratio`. `txn to pss ratio` determines the total number of transaction descriptors available to the server. At start time, this ratio is multiplied by the `number of user connections` parameter to create the transaction descriptor pool:

```
# of transaction descriptors = number of user
connections * txn to pss ratio
```

The default `txn to pss ratio` value, 16, ensures compatibility with earlier versions of Adaptive Server. Prior to version 12.0, Adaptive Server allocated 16 transaction descriptors for each user connection. In version 12.0 and later, the number of simultaneous transactions is limited only by the number of transaction descriptors available in the server.

For example, to allocate 25 transaction descriptors for every user connection, use the command:

```
sp_configure 'txn to pss ratio', 25
```

You must restart Adaptive Server for this change to take effect.

Using Adaptive Server Transaction Coordination Services

This chapter describes how to configure and use Adaptive Server transaction coordination services.

Topic	Page
Overview of transaction coordination services	13
Requirements and behavior	15
Configuring participant server resources	16
Using transaction coordination services in heterogeneous environments	18
Monitoring coordinated transactions and participants	19

Overview of transaction coordination services

The work of a local Adaptive Server transaction is sometimes distributed to remote servers that modify remote data. This can happen when a local transaction executes a remote procedure call (RPC) to update data in another Adaptive Server table, or when a local transaction modifies data in a remote table using Component Integration Services (CIS).

Prior to Adaptive Server version 12.0, local transactions that executed RPCs or updated data via CIS could not roll back the work done in remote Adaptive Servers. Moreover, the client executing the local transaction could not ensure that the remote work was actually committed if, for example, the remote server experienced a system failure.

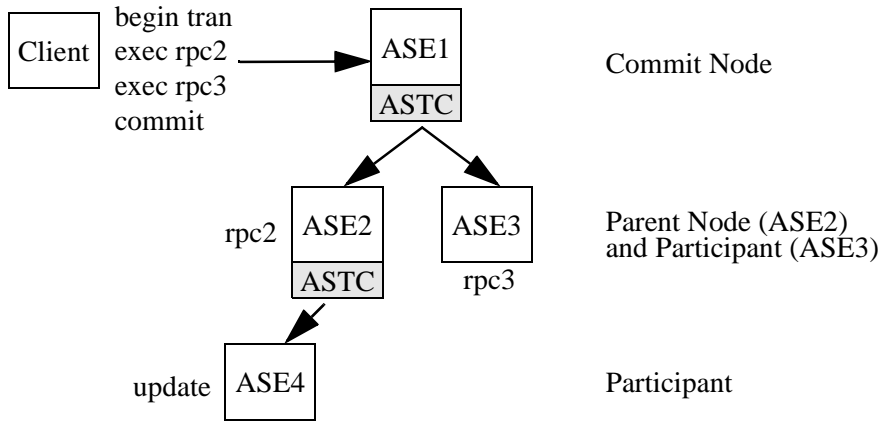
Adaptive Server provides services to propagate transactions to remote servers and coordinate the work of all servers, ensuring that all work is either committed or rolled back as a logical unit. With these transaction coordination services, Adaptive Server itself can act as a distributed transaction manager for transactions that update data in multiple Adaptive Servers.

Hierarchical transaction coordination

Because other Adaptive Servers involved in a distributed transaction may also coordinate remote participants, transactions can be further propagated to additional servers in a hierarchical manner. For example, in Figure 3-1, the client connected to ASE1 begins a transaction that executes an RPC on ASE2 and an RPC on ASE3. The coordination service for ASE1 propagates the transaction to ASE2 and ASE3.

Since ASE2 also has transaction coordination services enabled, it can propagate the transaction to additional remote participants. Here, ASE2 propagates the transaction to ASE4 where data is updated using CIS.

Figure 3-1: Hierarchical transaction coordination



In Figure 3-1, ASE1 is referred to as the commit node for the distributed transaction. When the transaction on ASE1 commits, the coordination service for ASE1 instructs ASE2 and ASE3 to prepare the transactions that it propagated to them. ASE3 indicates that its transaction is prepared when its local work is ready to be committed. ASE2 must complete its local work and instruct ASE3 to prepare its transaction. When the transactions are prepared in ASE2 and ASE3, the coordination service in ASE1 commits the original transaction. The instruction to commit subordinate transactions is then transmitted to ASE2, ASE3, and ultimately to ASE3, in the same manner as the instruction to prepare was transmitted.

X/Open XA-compliant behavior in DTP environments

The X/Open XA protocol requires resource managers to provide coordination services for transactions that are propagated to remote resource managers. This requirement is made because the external transaction manager (and in some cases, the client originating the transaction) has no knowledge of when transactions are propagated to remote servers, and therefore cannot ensure that the remote transactions complete or abort as required.

The new transaction coordination service brings Adaptive Server, in its role as a resource manager, into full compliance with the X/Open XA protocol. Distributed transactions can be implicitly propagated to remote servers through RPCs and CIS, and Adaptive Server guarantees that the commit or rollback status of the global transaction is preserved in the remote servers it coordinates.

Requirements and behavior

Adaptive Server transaction coordination services can ensure that the work of remote servers is logically committed or rolled back provided that each remote Adaptive Server is at version 12.0 and later.

Transaction coordination services are transparent to the client executing the distributed transaction. When a local client transaction executes a RPC or updates data via CIS, the coordination service creates a new transaction name for the remote work and propagates that transaction to the subordinate, remote server. When the local client commits or rolls back the local transaction, Adaptive Server coordinates that request with each of the subordinate servers to ensure that the remote transactions are committed or rolled back as well.

The Adaptive Server transaction coordination service runs as one or more background tasks named “ASTC HANDLER,” and can be viewed using `sp_who`. In systems using multiple Adaptive Server engines, the number of “ASTC HANDLER” processes (rounded down to the nearest whole number) is:

$$\text{number of engines} * 2/3$$

There can be a maximum of 4 “ASTC HANDLER” processes running on Adaptive Server.

The following output from `sp_who` shows a single “ASTC HANDLER”:

```
sp_who
fid spid status loginame origname  hostname blk_spid dbname cmd block_xlroid
```

```

-----
0  1  running  sa      sa      dtmsoll  0      master SELECT      0
0  2  sleeping NULL    NULL    master   0      master NETWORK HANDLER 0
0  3  sleeping NULL    NULL             0      master DEADLOCK TUNE   0
0  4  sleeping NULL    NULL             0      master MIRROR HANDLER  0
0  5  sleeping NULL    NULL             0      master HOUSEKEEPER     0
0  6  sleeping NULL    NULL             0      master CHECKPOINT SLEEP 0
0  7  sleeping NULL    NULL    metin1_dtm 0      sybsystemdb ASTC HANDLER 0

```

Configuring participant server resources

By default, the transaction coordination service is always enabled. The System Administrator can enable or disable these services using the `enable xact coordination` configuration parameter. See the *System Administration Guide* for a complete description of this parameter.

The System Administrator must also ensure that Adaptive Server has the required resources to coordinate all of the RPCs and CIS updates that may be requested by transactions. Each time a transaction issues an RPC or CIS update, the transaction coordinator must obtain a free **DTX participant**. A DTX participant or “distributed transaction participant” is an internal memory structure that Adaptive Server uses to coordinate a transaction that has been propagated to a subordinate Adaptive Server. In Figure 3-1 ASE1 requires three free DTX participants, and ASE2 requires two free DTX participants. (In each case, a single DTX participant is used to coordinate the local work of the transaction that is propagated.)

DTX participant resources remain in use by the coordinating Adaptive Server until the associated remote transaction has committed. This generally occurs some period of time *after* the initiating transaction has committed, since the initiating transaction commits as soon as all subordinate transactions have successfully prepared their work.

If no DTX participants are available, RPC requests and CIS update requests cannot proceed and the transaction is aborted.

***number of dtx participants* parameter**

The System Administrator can configure the total number of DTX participants available in Adaptive Server using the `number of dtx participants` configuration parameter. `number of dtx participants` sets the total number of remote transactions that the Adaptive Server transaction coordination service can propagate and coordinate at one time.

By default, Adaptive Server can coordinate 500 remote transactions. Setting `number of dtx participants` to a smaller number reduces the number of remote transactions that the server can manage. If no DTX participants are available, new distributed transactions will be unable to start. In-progress distributed transactions may abort if no DTX participants are available to propagate a new remote transaction.

Setting `number of dtx participants` to a larger number increases the number of remote transaction branches that Adaptive Server can handle, but also consumes more memory.

Optimizing *number of dtx participants* for your system

During a peak period, use `sp_monitorconfig` to examine the use of DTX participants:

```

                sp_monitorconfig "number of dtx participants"
Usage information at date and time: Jun 18 1999 9:00AM.
Name           # Free    # Active % Active  # Max Ever Used  Re-used
-----
number of dtx  480      20      4.00     210              N/A
participants

```

If the `#Free` value is zero or very low, new distributed transactions may be unable to start due to a lack of DTX participants. Consider increasing the `number of dtx participants` value.

If the `#Max Ever Used` value is too low, unused DTX participants may be consuming memory that could be used by other server functions. Consider reducing the value of `number of dtx participants`.

Using transaction coordination services in heterogeneous environments

When Adaptive Server propagates transactions to other version 12.0 and later Adaptive Servers, it can ensure the integrity of the distributed transaction as a whole. However, the work of a local Adaptive Server transaction is sometimes distributed to remote servers that do not support version 12.0 and later transaction coordination services. This may occur when a transaction uses RPCs to update data in earlier Adaptive Server versions, or when CIS services are used to update data in non-Sybase databases. Under these circumstances the coordinating Adaptive Server cannot ensure that the work of remote servers is committed or rolled back with the original transaction.

strict dtm enforcement parameter

In Adaptive Server, the System Administrator can enforce or relax the requirement to have distributed transactions commit or roll back as a logical unit by setting the `strict dtm enforcement` configuration parameter.

Note You can also override the value of `strict dtm enforcement` using the session level `set` command with the `strict_dtm_enforcement` option.

`strict dtm enforcement` determines whether or not Adaptive Server transaction coordination services will strictly enforce the ACID properties of distributed transactions.

Setting `strict dtm enforcement` to 1 (on) ensures that transactions are propagated only to servers that can participate in Adaptive Server-coordinated transactions. If a transaction attempts to update data in a server that does not support transaction coordination services, Adaptive Server aborts the transaction.

In heterogeneous environments, you may want to make use of servers that do not support transaction coordination. This includes older versions of Adaptive Server and non-Sybase database stores configured using CIS. Under these circumstances, you can set `strict dtm enforcement` to 0 (off). This allows Adaptive Server to propagate transactions to legacy Adaptive Servers and other data stores, but does not ensure that the remote work of these servers is rolled back or committed with the original transaction.

Monitoring coordinated transactions and participants

Adaptive Server tracks information about the status of work done in subordinate servers using data in the new system table, `sybsystemdbdbo.syscoordinations`. See the *Reference Manual* for a complete definition of this table.

The `sp_transactions` procedure also displays some data from the `syscoordinations` table for in-progress, remote transactions. See “Getting information about distributed transactions” on page 23 for more information.

DTM Administration and Troubleshooting

This chapter provides information about how to monitor, administer, and troubleshoot Adaptive Server DTM features.

Topic	Page
Transactions and threads of control	21
Getting information about distributed transactions	23
Crash recovery procedures for distributed transactions	29
Heuristically completing transactions	30

Transactions and threads of control

Prior to Adaptive Server version 12.0, all of a transaction's resources were privately owned by a single server task. The server could not share a transaction with any task other than the one that initiated the transaction.

Adaptive Server version 12.5 and later provides native support for the "suspend" and "join" semantics used by X/Open XA-compliant transaction managers such as Encina and TUXEDO. Transactions may be shared among different threads of execution, or may have no associated thread at all.

When a transaction has no thread associated with it, it is said to be "detached". Detached transactions are assigned a spid value 0. You can see the transaction spid value in the new `master.dbo.systransactions` table, or in output from the new `sp_transactions` procedure. See "Getting information about distributed transactions" on page 23 for more information.

Implications for system administrators

Detached transactions are meant to persist in Adaptive Server, since the client application may want to reattach the original thread, or attach a new thread to the transaction. The System Administrator can no longer roll back a transaction by killing its associated `spid`, as a thread is not attached to the transaction.

Transactions in a detached state may also prevent the log from being truncated with the `dump transaction` command. In extreme circumstances, detached transactions can be rolled back by using the new `dbcc complete_xact` command to heuristically complete a transaction. See “Heuristically completing transactions” on page 30.

dtm detach timeout period parameter

The system administrator can also specify a server-wide interval after which Adaptive Server automatically rolls back transactions that are in the detached state. `dtm detach timeout period` sets the amount of time, in minutes, that a distributed transaction branch can remain in the detached state. After this time has passed, Adaptive Server rolls back the detached transaction.

For example, to automatically rollback detached after 30 minutes, use the command:

```
sp_configure 'dtm detach timeout period', 30
```

Lock manager changes to support detached transactions

Prior to Adaptive Server version 12.0, the lock manager could uniquely identify a transaction’s locks by using the `spid` value of the transaction’s thread. With the new transaction manager, transactions may be detached from their original threads, and have no associated `spid`. Moreover, multiple threads with different `spid` values must be able to share the same transaction locks to perform the work of a distributed transaction.

To facilitate these changes, the Adaptive Server version 12.5 and later lock manager uses a unique lock owner ID, rather than a `spid`, to identify transaction locks. The lock owner ID is independent from the `spid` that created the transaction, and it persists even when the transaction is detached from a thread. Lock owner IDs provide a way to support transactional locks when transactions have no associated threads, or when a new thread is attached to the transaction.

The lock owner ID is stored in the new `loid` column of `master.dbo.syslocks`. You can determine the `loid` value of a transaction by examining `sp_lock` or `sp_transactions` output.

Examining the `spid` and `loid` columns from `sp_transactions` output provides information about a transaction and its thread of control. A `spid` value of zero indicates that the transaction is detached from its thread of control. Non-zero `spid` values indicate that the thread of control is currently attached to the transaction.

If the `loid` value in `sp_transactions` output is even, then a local transaction owns the lock. Odd `loid` values indicate that an external transaction owns the lock.

See “Getting information about distributed transactions” on page 23 for more information about `sp_transactions` output.

Getting information about distributed transactions

Adaptive Server has system table, `master.dbo.systransactions`, which stores information about all server transactions. `systransactions` identifies each transaction and maintains information about the state of the transaction and its associated threads.

The new system procedure, `sp_transactions`, translates information from the `systransactions` and `syscoordinations` tables to display status conditions for active transactions.

Transaction identification in *systransactions*

Adaptive Server stores transaction names in a column of `varchar(255)` (as compared to `varchar(64)` in previous server versions) to accommodate the length and format of transaction names supplied by different distributed transaction protocols. In the X/Open XA protocol, for instance, distributed transactions are assigned a transaction name consisting of both a global transaction ID (`gtrid`) and a branch qualifier. Within Adaptive Server, this information is combined in the `xactname` column of the `systransactions` table.

`systransactions.xactname` stores the names of both externally-created distributed transactions (defined by an X/Open XA transaction manager or MSDTC) and local server transactions. Clients defining local transactions can name those transactions anything they wish, within the confines of the `varchar(255)` column. Similarly, external transaction managers can use a variety of different formats to name a distributed transaction.

Transaction keys

The transaction key, stored in the `xactkey` column of `systransactions`, acts as a unique internal handle to a server transaction. For local transactions, `xactkey` ensures that transactions can be distinguished from one another, even if the transaction name is not unique to the server.

Beginning with Adaptive Server version 12.0, all system tables refer to `systransactions.xactkey` to uniquely identify a transaction. The `sysprocesses` and `syslogshold` tables are the only exceptions to this rule—they reference `systransactions.xactname` and truncate the value to a length of `varchar(64)` (for `sysprocesses`) and `varchar(67)` (for `syslogshold`), to maintain backward compatibility with earlier Adaptive Server versions.

Viewing active transactions with `sp_transactions`

The `sp_transactions` procedure translates information from `systransactions` and `syscoordinations` to provide information about active transactions. When used without keywords, `sp_transactions` displays information about all active transactions:

```

                sp_transactions
xactkey          type      coordinator starttime
state           connection dbid  spid   loid
failover                srvname                namelen
xactname
-----
-----
-----
-----
0x00000b1700040000dd6821390001 Local      None      Jun 1 1999 3:47PM
Begun           Attached    1    1      2
Resident Tx                NULL
$user_transaction
0x00000b1700040000dd6821390001 Remote     ASTC      Jun 1 1999 3:47PM

```

```

Begun          NA          0  8  0
Resident Tx    caserv2          108

00000b1700040000dd6821390001-aa01f04ebb9a-00000b1700040000dd6821390001-
aa01f04ebb9a-caserv1-caserv1-0002

```

Identifying local, remote, and external transactions

The “type” column indicates whether the transaction is local, remote, or external. Local transactions execute on the local server (the server on which you ran `sp_transactions`). Local transactions have a null value in the “srvname” column, since the transaction takes place on the current server.

For remote transactions, `sp_transactions` lists the name of the server executing the transaction under the “srvname” column. The `sp_transactions` output above shows a remote transaction executing on the server named `caserv2`.

External transactions indicate that the transaction is coordinated by an external transaction coordinator, such as CICS, Encina, or the “ASTC HANDLER” process of another Adaptive Server. External transactions also have a null value in the “srvname” column.

Identifying the transaction coordinator

The “coordinator” column indicates the method or protocol used to manage a transaction. In the output above, the local transaction `$user_transaction` does not have an external coordinator. The remote transaction taking place on `caserv2` has the coordinator value “ASTC”. This indicates that the transaction is coordinated using native Adaptive Server coordination services, as described under Chapter 3, “Using Adaptive Server Transaction Coordination Services”.

See `sp_transactions` in the *Reference Manual* for a complete list and description of possible coordinator values.

Viewing the transaction thread of control

The `spid` column displays the Process ID of the process attached to the transaction (or 0 if the transaction is detached from its thread of control). For local transactions, the `spid` value indicates a Process ID running on the local server. For remote transactions, the `spid` indicates the Process ID of a task running on the indicated remote server. The output above shows a `spid` value of 8 running on the remote server, `caserv2`.

Understanding transaction state information

The “state” column displays information about the current state of each transaction. At any given time, a local or external transaction may be executing a command, aborted, committed, and so forth. Additionally, distributed transactions can be in a prepared state, or can be heuristically completed or rolled back.

The “connection” column displays information about the state of the transaction’s connection. You can use this information to determine whether a transaction is currently attached to or detached from a process. Transactions in X/Open XA environments may become detached from their initiating process, in response to requests from the transaction manager.

See `sp_transactions` in the *Reference Manual* for a complete list and description of possible coordinator values.

Limiting `sp_transactions` output to specific states

You can use `sp_transactions` with the `state` keyword to limit output to the specified transaction state. For example:

```
sp_transactions "state", "Prepared"
```

displays information only for distributed transactions that have been prepared.

Transaction failover information

The “failover” column displays special information for servers operating in high availability environments. In high availability environments, prepared transactions may be transferred to a secondary companion server if the original server experiences a critical failure. The “failover” column can display three possible failover states that indicate how and where the transaction is executing:

- “Resident Tx” is displayed under normal operating conditions, and on systems that do not utilize Adaptive Server high availability features. “Resident Tx” means that the transaction was started and is executing on a primary Adaptive Server.


```
Resident Tx                NULL                108

00000b1700040000dd6821390001-aa01f04ebb9a-00000b1700040000dd6821390001-
aa01f04ebb9a-caserv1-caserv1-0002

caserv1 caserv1
00000b1700040000dd6821390001-aa01f04ebb9a
```

Commit and parent nodes

For distributed transactions coordinated by Adaptive Server, the “commit node” column lists the name of the server that executes the topmost branch of the distributed transaction. This transaction determines the commit or rollback status for all branches of the transaction. See “Hierarchical transaction coordination” on page 14 for more information.

The “parent node” column lists the name of the server that initiated the transaction. In the `sp_transactions` output above, the “commit node” and “parent node” columns list the same server, `caserv1`. This indicates that the distributed transaction originated on `caserv1`, and `caserv1` propagated a branch of the transaction to the current server.

Global transaction ID

The “gtrid” column displays the global transaction ID for distributed transactions coordinated by Adaptive Server. Transaction branches that are part of the same distributed transaction share the same gtrid. You can use a specific gtrid with the `sp_transactions` `gtrid` keyword to determine the state of other transaction branches running on the current server. This is useful for System Administrators who must determine whether a particular branch of a distributed transaction should be heuristically committed or rolled back. See “Determining the commit status for Adaptive Server transactions” on page 33 for an example that uses `sp_transactions` with the `gtrid` keyword.

Note For transactions coordinated by an X/Open XA-compliant transaction manager, MSDTC, or SYB2PC, the `gtrid` column shows the full transaction name supplied by the external coordinator.

Crash recovery procedures for distributed transactions

During crash recovery, Adaptive Server must resolve distributed transactions that it discovers in the prepared state. The method used to resolve prepared transactions depends on the coordination method or coordination protocol used to manage the distributed transaction.

Note The following crash recovery procedures are not performed during normal database recovery for `load database` or `load transaction` commands. If `load database` or `load transaction` applies any transactions that prepared or in-doubt, Adaptive Server aborts those transactions before bringing the associated database online.

Transactions coordinated with MSDTC

Prepared transactions that were coordinating using MSDTC are rolled forward or backward depending on the commit status of the master transaction. During recovery, Adaptive Server initiates contact with MSDTC to determine the commit status of the master transaction, and commits or rolls back the prepared transaction accordingly. If it cannot contact MSDTC, the recovery procedure waits until contact is established. Further recovery does not take place until Adaptive Server has established contact with MSDTC.

Transactions coordinated by Adaptive Server or X/Open XA

During crash recovery, Adaptive Server may also encounter prepared transactions that were coordinated using Adaptive Server transaction coordination services or the X/Open XA protocol. Upon encountering these transactions, the local server must wait for the coordinating Adaptive Server or the external transaction coordinator to initiate contact and indicate whether the prepared transaction should commit or roll back.

To speed the recovery process, Adaptive Server restores each of these transactions to their condition prior to the failure. The transaction manager creates a new transaction with the original transaction ID, and the lock manager applies locks to protect data that the original transaction was modifying. The restored transaction remains in a prepared state but is detached, having no thread associated with it.

Once the transaction's coordinator contacts Adaptive Server, the transaction manager can commit or roll back the transaction.

Using this recovery mechanism, the server can bring a database online even when the coordinating Adaptive Server or external transaction manager has not yet attempted to resolve the prepared transaction. Other clients and transactions can resume work on the local data, since the prepared transaction holds the locks it did prior to recovery. The prepared transaction itself is ready to commit or roll back once contacted by its coordinator.

When the controlling Adaptive Server or external transaction manager cannot complete the transaction, the System Administrator can heuristically complete the transaction to free its locks and transaction resources. See "Heuristically completing transactions" on page 30 for more information.

Transactions coordinated with SYB2PC

Prepared transactions that were coordinated using the SYB2PC protocol are rolled forward or backward depending on the commit status of the master transaction. During recovery, Adaptive Server initiates contact with the commit service to determine the commit status of the master transaction, and commits or rolls back the prepared transaction accordingly. If it cannot contact the commit service, Adaptive Server does not bring the database online. However, Adaptive Server does proceed to recover other databases in the system.

This recovery method was used for SYB2PC transactions in earlier Adaptive Server versions and is unchanged with Adaptive Server version 12.5 and later.

Heuristically completing transactions

Adaptive Server includes the `dbcc complete_xact` command to facilitate heuristic completion of transactions. `dbcc complete_xact` resolves a transaction by either committing or rolling back its work, freeing whatever resources the transaction was using.

`dbcc complete_xact` is provided for those cases where only the System Administrator can properly resolve a prepared transaction, or for when the System Administrator must resolve a transaction without waiting for the transaction's coordinator.

For example, in Figure 3-1 on page 14, heuristic completion may be considered if all remote Adaptive Servers have prepared their transactions, but the network connection to ASE1 was *permanently lost*. The remote Adaptive Servers will maintain their transactions in a prepared state until contacted by the coordination service from ASE1. In this case, only the System Administrator for ASE2, ASE3, and ASE4 can properly resolve the prepared transactions. Heuristically completing the prepared transaction in ASE3 frees up transaction and lock resources, and records the commit status in `systransactions` for later use by the transaction coordinator. Heuristically completing the transaction in ASE2 also completes the transaction propagated to ASE4.

Completing prepared transactions

Warning! Heuristically completing a prepared transaction can cause inconsistent results for an entire distributed transaction. The System Administrator's decision to heuristically commit or roll back a transaction may contradict the decision made by the coordinating Adaptive Server or transaction protocol.

Before heuristically completing a transaction, the System Administrator should make every effort to determine whether the coordinating Adaptive Server or transaction protocol decided to commit or roll back the distributed transaction (see “Determining the commit status for Adaptive Server transactions” on page 33).

By using `dbcc complete_xact`, the System Administrator forces Adaptive Server to commit or roll back a branch of a distributed transaction. After heuristically completing a prepared transaction, Adaptive Server records the transaction's commit status in `master.dbo.systransactions` so that the transaction's coordinator—Adaptive Server, MSDTC, or an X/Open XA transaction manager—can know whether the transaction was committed or rolled back.

Adaptive Server propagates the command to heuristically commit or abort a transaction to any participant servers that it coordinated for the transaction branch. For example, if in Figure 3-1 on page 14 you heuristically commit the transaction on ASE2, ASE2 propagates the command to ASE4 so that the transaction on ASE4 also commits.

`dbcc complete_xact` requires that you supply an active transaction name and desired outcome for the transaction. For example, the following command heuristically commits a transaction:

```
dbcc complete_xact "00000b1700040000dd6821390001-  
aa01f04ebb9a-00000b1700040000dd6821390001-  
aa01f04ebb9a-caserv1-caserv1-0002", "commit"
```

Forgetting heuristically completed transactions

When the System Administrator heuristically completes a prepared transaction, Adaptive Server maintains information about the transaction's commit status in `master.dbo.systransactions`. This information is maintained so external transaction coordinators can detect the presence of heuristically completed transactions.

If the external coordinator is another Adaptive Server, the server examines the commit status and logs a warning message if the heuristic completion conflicts with the commit status of the distributed transaction. After examining the commit status, the coordinating Adaptive Server clears the commit status information from `systransactions`.

If the external coordinator is an X/Open XA-compliant transaction manager, the transaction manager does not log warning message when the heuristic completion conflicts with the distributed transaction. However, X/Open XA-compliant transaction managers clear the commit status information from `systransactions`.

Manually clearing the commit status

`dbcc forget_xact` purges the commit status of a heuristically completed transaction from `systransactions`. It can be used when the System Administrator does not want the coordinating service to have knowledge that a transaction was heuristically completed, or when an external coordinator will not be available to clear information from `systransactions`.

See `dbcc` in the *Reference Manual* for more information about using `dbcc forget_xact`.

Completing transactions that are not prepared

`dbcc complete_xact` can also be used to roll back Adaptive Server-coordinated transactions that have not yet reached the prepared state. Heuristically rolling back a transaction that has not yet been prepared does not pose a risk to the distributed transaction, since the coordinating server can recognize that the transaction failed to prepare its work. Under these circumstances, the coordinating Adaptive Server can roll back the entire distributed transaction to preserve consistency.

When you heuristically roll back an Adaptive Server transaction that has not yet been prepared, Adaptive Server *does not* record the heuristic roll back in `systransactions`. Instead, an informational message is printed to the screen and recorded in the server's error log.

Determining the commit status for Adaptive Server transactions

If the distributed transaction branch you want to commit or roll back is coordinated by Adaptive Server, you can use `sp_transactions` to determine the commit status of the distributed transaction. To do so, complete the following steps.

Note These steps cannot be used with distributed transactions that are coordinated by the X/Open XA protocol, MSDTC, or SYB2PC.

- 1 In the server that is executing the transaction branch you want to complete, use `sp_transactions` with the `xid` keyword to display information about the transaction. Record the commit node and `gtrid` of the transaction. For example:

```

sp_transactions "xid",
"00000b1700040000dd6821390001-aa01f04ebb9a-
00000b1700040000dd6821390001-aa01f04ebb9a-caserv1-
caserv1-0002"
xactkey                               type      coordinator starttime
state                                 connection dbid   spid   loid
failover                               srvname
xactname
commit_node parent_node
gtrid
-----
-----
-----

```

```

-----
-----
-----
-----
0x00000b2500080000dd6821960001 External ASTC Jun 1 1999 3:47PM
Begun Attached 1 8 139
Resident Tx NULL 108
    
```

```

00000b1700040000dd6821390001-aa01f04ebb9a-
00000b1700040000dd6821390001-aa01f04ebb9a-caserv1-caserv1-0002

caserv1 sfserv
00000b1700040000dd6821390001-aa01f04ebb9a
    
```

In this example, the commit node for the distributed transaction is “caserv1” and the gtrid is “00000b1700040000dd6821390001-aa01f04ebb9a”.

- 2 Log on to the server indicated by the commit node. For example:


```
isql -Usa -Psa_password -Scaserv1
```
- 3 Use sp_transactions with the gtrid keyword to determine the commit status of the distributed transaction having the gtrid obtained in step 1:


```
sp_transactions "gtrid",
                "00000b1700040000dd6821390001-aa01f04ebb9a"
```

```

xactkey type coordinator starttime
state connection dbid spid loid
failover srvname namelen
xactname
commit_node
parent_node
    
```

```

-----
-----
-----
-----
-----
-----
-----
-----
0x00000b1700040000dd6821390001 Local None Jun 1 1999 3:47PM
Committed Attached 1 1 2
Resident Tx NULL 17
$user_transaction

caserv1
caserv1
    
```

In this example, the local transaction with the specified gtrid has committed, as indicated by the “state” column. The System Administrator should heuristically *commit* the prepared transaction examined in step 1.

- 4 Using an account with System Administrator privileges, log on to the server that is executing the transaction branch you want to complete:

```
isql -Usa -Psa_password -Ssfserv
```

- 5 Use `dbcc complete_xact` to commit the transaction. In this example, the System Administrator should use the `commit` keyword to maintain consistency with the distributed transaction:

```
dbcc complete_xact "00000b1700040000dd6821390001-  
aa01f04ebb9a-00000b1700040000dd6821390001-  
aa01f04ebb9a-caserv1-caserv1-0002", "commit"
```

Configuration considerations

This section describes configuration options to consider when troubleshooting.

Behavior of DDLs within distributed transactions

If a transaction is coordinated by an external transaction manager using X/Open XA protocol or through Adaptive Server transaction coordination services of another Adaptive Server, then DDL commands are not allowed within the transaction. This behavior applies even if the database option `ddl in tran` is enabled.

Index

Symbols

::= (BNF notation)

in SQL statements x

, (comma)

in SQL statements x

{ } (curly braces)

in SQL statements x

() (parentheses)

in SQL statements x

[] (square brackets)

in SQL statements x

A

accessing DTM 7

asset management 7

ASTC Handler 15

B

Backus Naur Form (BNF) notation x

begin transaction 9

behavior

CIS transactions 4

coordination services 15

manager-coordinated transaction 3

RPC transactions 4

x/Open XA-compliant 15

BNF notation in SQL statements x

brackets. *See* square brackets []

C

case sensitivity

in SQL xi

CICS 3, 25

CIS 1, 4, 13, 15

comma (,)

in SQL statements x

commands

begin transaction 9

dbcc 22, 30, 31, 32

dbcc complete_xact 22, 30, 31

dbcc forget_xact 32

dump transaction 22

load database 29

load transaction 29

update 11

commit node 27, 28

commit status 33

complete_xact 22, 30, 31

Component Integration Services

See CIS

configuration parameters

See parameters

configuring

DTX participants 17

participant server resources 16

transaction resources 8

conventions

See also syntax

Transact-SQL syntax x

used in the Reference Manual x

coordinated transactions 19

coordination services 3, 13–19

behavior of 15

CIS and 13

heterogeneous environments and 18

- hierarchical 14
- overview of 2
- requirements of 15
- RPCs and 13
- crash recovery 29
- curly braces ({ }) in SQL statements x
- D
- dbcc 2, 22, 30, 31, 32
- detached transactions 21
 - spid value and 21
- Distributed Transaction Management
 - See DTM
- DTM
 - accessing 7
 - administration of 21–35
 - coordination services 2
 - enabling 7
 - features 1
 - licensing of 7
 - overview of 1–5
 - prepared transaction recovery 2
 - starting 7
 - thread management with 2
 - transaction descriptors 9
 - troubleshooting of 21–35
 - XA interface and 3
- dtm detach timeout period 22
- dtm_tm_role 3
- DTX participant 16
 - configuring 17
 - optimizing 17
- dump transaction 22
- E
- enable dtm 8
- enable xact coordination 8, 16
- Encina 3
- external transactions 25
- F
- failed-over Tx status 27
- failover information 26
- failover state
 - failed-over Tx status 27
 - Resident Tx 26
 - Tx by Failover-Conn 27
- forget_xact 32
- G
- global transaction ID (gtrid) 27, 28
- H
- heuristic completion 22, 30–35
 - forgetting 32
- high availability 26
- I
- installation
 - of license keys 7
- isql 11
- L
- license key 7
- load database 29
- load transaction 29
- local transactions 25
- lock manager 22
- loid 22
 - even value 23
 - odd value 23
- M
- Microsoft Distributed Transaction Coordinator
 - See MSDTC
- monitoring 19
 - participants 19
- MSDTC 2, 3, 8, 24, 29
 - dtm_tm_role and 3

- ODBC driver and 3
- XA interface and 3
- multi-database transactions 9
- N
- nodes
 - commit 28
 - parent 28
- number of dtx participants 17
 - optimizing 17
- number of engines 15
- number of user connections 11
- O
- ODBC driver 3
- P
- parameters
 - dtm detach timeout period 22
 - enable dtm 8
 - enable xact coordination 8
 - number of dtx participants 17
 - number of engines 15
 - number of user connections 11
 - strict dtm enforcement 18
 - txn to pss ratio 9, 11
- parent node 28
- parentheses ()
 - in SQL statements x
- participants
 - See DTX participant
- prepared transactions 33
- R
- recovering transactions 2
- remote procedure call
 - See RPCs
- remote transactions 25
- resident Tx 26
- resources
 - participant server 16
 - transaction descriptor 9
- RPCs 4, 13, 15
- S
- shared transactions 21
- sp_addobjectdef 4
- sp_configure 7, 11
- sp_transactions 19, 21, 23, 23–28
- sp_who 15
- spid 21
 - lock manager and 22
 - multiple threads and 22
 - Process ID and 25
 - roll back of transactions and 22
 - thread of control and 23
 - transaction manager 22
- square brackets []
 - in SQL statements x
- srvname column 25
- starting DTM 7
- state information 26
- stored procedures
 - sp_addobjectdef 4
 - sp_configure 7, 11
 - sp_lock 23
 - sp_monitorconfig 17
 - sp_transaction 23–28
 - sp_transactions 19, 21, 23, 33
 - sp_who 15
- strict dtm enforcement 18
- SYB2PC transactions 5, 30
- sybssystemdb database 19
- symbols
 - in SQL statements x
- syntax conventions, Transact-SQL x
- SySAM 7

- syscoordinations table 19, 23, 24
- syslocks table 23
- syslogshold table 24
- sysprocesses table 24
- systransactions table 21, 23, 31, 32
- T
- tables
 - syscoordinations 19, 23, 24
 - syslocks 23
 - syslogshold 24
 - sysprocesses 24
 - systransactions 21, 23, 31, 32
- thread of control 25
- threads 21
- transaction descriptors 9, 11
- transactions
 - Adaptive Server coordination of 3, 13–19, 25
 - CIS and 4
 - detached 21
 - determining commit status of 33
 - external 24, 25
 - failover status of 26
 - gtrid of 27, 28
 - hierarchical coordination of 14
 - initiating 9
 - keys of 24
 - local 25
 - manager-coordinated transaction behavior 3
 - monitoring 19
 - MSDTC coordination of 29
 - multi-database 9
 - outer level 9
 - prepared 33
 - recovery of 2, 29
 - remote 25
 - resources for 8
 - RPCs and 4
 - sharing 21
 - spid value and 21
 - state information of 26
 - SYB2PC coordination of 5, 30
 - threads 21, 25
 - TM coordination of 3
 - types of 2
 - X/Open XA coordination of 29
- TUXEDO 3
- Tx by Failover-Conn 27
- txn to pss ratio 9, 11
- U
- update 11
- V
- varchar(255) 23, 24
- varchar(64) 23, 24
- varchar(67) 24
- X
- X/Open XA 3, 8, 15, 24, 29
- XA interface 3
- xactkey column 24
- xactname column 23
- XA-Server 1, 3
- XA-Server product 3
- xid 27