



# **Migrating to Adaptive Server Enterprise 12.5**

**Adaptive Server Enterprise  
12.5**

DOCUMENT ID: 34982-01-1250-01

LAST REVISED: Oct 2001

Copyright © 1989-2001 by Sybase, Inc. All rights reserved.

This publication pertains to Sybase database management 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 Server IQ, Adaptive Warehouse, AnswerBase, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-FORMS, APT-Translator, APT-Library, Backup Server, ClearConnect, Client-Library, Client Services, Convoy/DM, Copernicus, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DB-Library, dbQueue, Developers Workbench, Direct Connect Anywhere, DirectConnect, Distribution Director, e-ADK, E-Anywhere, 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, ImpactNow, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, KnowledgeBase, MainframeConnect, Maintenance Express, MAP, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, MySupport, Net-Gateway, Net-Library, New Era of Networks, ObjectConnect, ObjectCycle, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Biz, Open Client, Open ClientConnect, 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, Power++, power.stop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, PowerJ, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Rapport, Report Workbench, Report-Execute, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Resource Manager, RW-DisplayLib, RW-Library, S-Designor, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILLS, 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 MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, Transact-SQL, Translation Toolkit, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, Viewer, Visual Components, VisualSpeller, 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 and XP Server are trademarks of Sybase, Inc. 8/01

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., 6475 Christie Avenue, Emeryville, CA 94608.

- Introduction and Guide to Resources 1
- In This Chapter 1
- Chapters in the Migration Guide 1
- Relating Documentation to Migration Phase 2
- Other Sources of Information 3
- Sybase Certifications on the Web 4
- Migration Web Page 5
- Classes 5
- Consulting Services 5
- If You Need Help 6
- Analyze: Documenting Business Requirements 1
- In This Chapter 1
- Diagramming the System 1
- Operational Business Requirements 3
- Availability Requirements 3
- Database Change Metrics 4
- Database Dump Details 4
- Maintenance Procedures 4
- Service Level Requirements 5
- Transaction Profile 5
- Current Performance Metrics 6
- Additional Business Requirements 7
- Analyze: Documenting Your Environment 1
- In This Chapter 1
- Hardware Configuration 1
- General Server Hardware 2
- CPU Resources per Machine 2
- Disk Configuration 2
- Network Configuration 4
- Tape Configuration 4
- Physical Memory Utilization 4
- Software Configuration 5
- Operating System 5
- Applications 6
- Sybase Configuration 6
- General Information 7

---

|   |    |
|---|----|
| Database Devices  | 7  |
| Databases and Segments  | 8  |
| Dump Devices  | 8  |
| Adaptive Server Objects   | 9  |
| Gather Scripts to Create Objects  | 9  |
| If You Don't Have Scripts   | 9  |
| Plan: Writing a Plan and Getting Ready to Migrate                       | 1  |
| In this Chapter   | 1  |
| Review Documentation  | 1  |
| Review the Upgrade Process  | 2  |
| Special Cases: Migrating to 64-bit Operating System or Larger Page Size | 3  |
| Determine Migration Approach  | 4  |
| Parallel With Replication   | 5  |
| Cutover Without Replication   | 7  |
| Phased Cutover  | 9  |
| Write a Migration Plan  | 11 |
| Build the Adaptive Server Environment                                   | 11 |
| Update Hardware Resources   | 12 |
| Verify Operating System Version and EBF Level                           | 12 |
| Review Adaptive Server Interoperability with Other Sybase Products      | 13 |
| Update Applications and System Administration Procedures                | 13 |
| Create Migration Scripts  | 14 |
| Implement: Making Required Application Changes                          | 1  |
| In This Chapter   | 1  |
| If Your Version is 11.5   | 2  |
| Changes to Locking  | 2  |
| Changes to Error Checking   | 4  |
| 11.9.2 Query Processing Changes   | 8  |
| If Your Version is 11.5 or 11.9.x                                       | 12 |
| ANSI Joins  | 12 |
| dsync Function  | 13 |
| 12.0 Query Processing Changes   | 14 |
| If Your Version is 11.5, 11.9.x, or 12.0                                | 17 |
| Changes to Transact-SQL   | 17 |
| The enable xact coordination Configuration Parameter                    | 17 |
| New Reserved Words  | 18 |

---

|   |    |
|---|----|
| Wide Columns and Data Truncation                  | 19 |
| Implement: Making Database Administration Changes | 1  |
| In This Chapter                                   | 1  |
| If Your Version is 11.5                           | 2  |
| New Rollback Records                              | 2  |
| Optimizer Changes                                 | 3  |
| Trace Flag 326                                    | 4  |
| If Your Version is 11.5 or 11.9.x                 | 4  |
| Documentation Changes                             | 5  |
| Backup Server                                     | 5  |
| Suspending Updates to Databases for Copying       | 7  |
| New Unix File System Device Support               | 8  |
| Changes to dbids                                  | 8  |
| CIS Proxy Databases                               | 9  |
| ASE Plug-in for Sybase Central                    | 9  |
| The sybsyntax Utility                             | 9  |
| If Your Version is 11.5, 11.9.x, or 12.0          | 9  |
| Changes to Documentation                          | 10 |
| System Database Sizes                             | 10 |
| Changes to System Tables                          | 11 |
| New Database and Table Limits                     | 12 |
| New Reserved Words                                | 13 |
| Configuration Parameters                          | 14 |
| Increased Memory                                  | 16 |
| New Storage for text and image Data               | 16 |
| buildmaster Command Obsolete                      | 17 |
| Test: Ensuring Stability and Performance          | 1  |
| In This Chapter                                   | 1  |
| The Goal of Testing                               | 1  |
| Setting Up the Test Environment                   | 2  |
| Make Backups                                      | 2  |
| Use Scripts to Create the Test System             | 2  |
| Create Your Databases by Loading Backups          | 3  |
| If the Test Environment Is Not an Exact Duplicate | 3  |
| Prioritizing Applications to be Tested            | 4  |
| Establishing Performance Criteria                 | 4  |

---

|  |    |
|--|----|
| Developing Fallback Procedures                     | 5  |
| Summary of Testing Techniques                      | 5  |
| Writing Performance Scripts                        | 7  |
| Write Benchmark Scripts                            | 7  |
| Drivers  | 8  |
| Test Cycle: Summary of Tests                       | 10 |
| Test Cycle: Testing for Performance                | 11 |
| Pre-Upgrade Single-User Tests                      | 11 |
| Pre-Upgrade Multi-User Tests                       | 12 |
| Test System Upgrade                                | 13 |
| Post-Upgrade Single-User Tests                     | 13 |
| Post-Upgrade Multi-User Tests                      | 13 |
| Worksheets for Your Current Environment            | 1  |
| Adaptive Server Operational Worksheets             | 1  |
| Operational Business Requirements                  | 1  |
| Backup and Restore Procedures                      | 2  |
| Database Dump Details                              | 3  |
| Maintenance Procedure Details                      | 4  |
| Data Architecture Worksheet                        | 5  |
| Client Application Components                      | 5  |
| Production Performance Metrics                     | 6  |
| Transaction Profile                                | 7  |
| Adaptive Server Infrastructure Worksheets          | 8  |
| Host Configuration                                 | 9  |
| Adaptive Server Configuration                      | 15 |
| Database Devices                                   | 16 |
| Databases and Segments                             | 17 |
| Dump Devices                                       | 18 |
| Sample Migration Task Lists                        | 1  |
| Sample Task List Template                          | 1  |
| General Migration Task List Example                | 2  |
| Migration Analysis                                 | 2  |
| Migration Preparation                              | 4  |
| Implement Migration (Using Install/Load Technique) | 5  |
| Implement Migration (Using Upgrade Technique)      | 7  |
| Migration Quality Assurance                        | 8  |

|  |    |
|--|----|
| Parallel Migration Task List Example 9                         |    |
| Define Test/Acceptance Criteria—Regression Test Suites         | 10 |
| Set Up Target Production Environment                           | 11 |
| Set Up Replication Server                                      | 11 |
| Run Regression Test Suites                                     | 11 |
| Upgrade Server B (Shadow)                                      | 12 |
| Run Post-upgrade Regression Test Suites on ASE 12.5 (Server B) | 13 |
| Run User Acceptance Tests on ASE 12.5 (Server B)               | 13 |
| Shift Production Users to Adaptive Server 12.5 (Server B)      | 14 |
| Perform Final Steps  | 14 |
| Cutover Migration Task List Example 14                         |    |
| Set Up ASE 12.5 Environment on Development System              | 15 |
| Define Test/Acceptance Criteria—Regression Test Suites         | 15 |
| Define Fallback Procedures on Test System                      | 16 |
| “Baseline” Older Environment on Test System                    | 16 |
| Run Regression Test Suites on Older Release Test System        | 16 |
| Upgrade Test System to Release 12.5                            | 17 |
| Run Regression Test Suites on Release 12.5 Test System         | 17 |
| Run User/Acceptance Tests on Release 12.5 Test System          | 18 |
| Execute Fallback Procedures on Test System                     | 18 |
| Upgrade Production Server to ASE 12.5                          | 18 |
| Perform Final Steps  | 19 |
| Staged Cutover Task Overview                                   | 19 |
| Tasks  | 19 |
| Migration Issue Checklists                                     | 1  |
| Logical Data Architecture                                      | 1  |
| Logical Application Architecture                               | 2  |
| Logical Technology Architecture                                | 2  |
| Logical Support Architecture                                   | 3  |
| Migration Strategy Design                                      | 4  |
| Pre-Upgrade Checklist  | 1  |
| Pre-Upgrade Checklist  | 1  |

---



# Introduction and Guide to Resources

## In This Chapter

This chapter gives an overview of the topics covered in this manual. In addition, it points you to a list of Sybase resources that can help you plan and execute a trouble-free migration to Sybase Adaptive Server Enterprise™ 12.5.

This chapter covers the following topics:

- Chapters in the Migration Guide
- Relating Documentation to Migration Phase
- Other Sources of Information
- Sybase Certifications on the Web
- Migration Web Page
- Classes
- Consulting Services
- If You Need Help

## Chapters in the Migration Guide

This manual is based on SAFE/EM migration methods. SAFE/EM, a methodology for migration developed by Sybase Professional Services, is organized into the following phases:

- *Analyze*
- *Prepare*
- *Implement*
- *Test*

The chapters in this manual are described in the following table:

| Chapter  | Contents   |
|--|--|
| Chapter 1, "Introduction and Guide to Resources"               | A survey of resources available from Sybase and third-party vendors.   |
| Chapter 2, "Analyze: Documenting Business Requirements"        | A series of worksheets for documenting the flow of information in your system and your business requirements.  |
| Chapter 3, "Analyze: Documenting Your Environment"             | A series of worksheets for documenting your environment, including hardware and software, for compatibility assessment.  |
| Chapter 4, "Plan: Writing a Plan and Getting Ready to Migrate" | Guidelines for determining the best migration plan for your site.  |
| Chapter 5, "Implement: Making Required Application Changes"    | A technical summary T-SQL syntax, query processing, reserved word and system changes that can affect your applications and cause failure and unexpected results. |
| Chapter 6, "Implement: Making Database Administration Changes" | A technical summary system changes that can require you to make adjustments to your database administration procedures.  |
| Chapter 7, "Test: Ensuring Stability and Performance"          | Guidelines for setting up a test environment and writing test suites.  |

## Relating Documentation to Migration Phase

Sybase provides documentation for all stages of the migration process. While this migration guide documents the minimum changes to your system and applications necessary to avoid problems, we recommend that you refer to What's New in Adaptive Server Enterprise at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic__BookView) and other Sybase manuals to help you plan the design of your new Adaptive Server system to take advantage of Sybase's new performance features.

The following table gives general guidelines for relating migration phase to Sybase documentation:

| Document  | Time period covered   | Range of tasks covered  |
|---|---|---|
| <ul style="list-style-type: none"> <li>What's New in Adaptive Server Enterprise at <a href="http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic__BookView">http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic__BookView</a></li> <li><i>Migrating to Adaptive Server Enterprise 12.5</i></li> </ul>  | Planning/preparation prior to upgrade                             | <ul style="list-style-type: none"> <li>Assessing current system</li> <li>Planning migration</li> <li>Making applications compatible</li> <li>Updating DBA procedures</li> </ul> |
| <ul style="list-style-type: none"> <li>Release Bulletins</li> <li>TechNotes (ASE Migration Resources Web page at <a href="http://sybase.com/support/techdocs/migration">http://sybase.com/support/techdocs/migration</a>)</li> </ul>  | Planning/preparation prior to upgrade                             | Finding information needed to avoid upgrade problems, including problem reports, special installation issues, and compatibility issues  |
| <ul style="list-style-type: none"> <li>installation guide at <a href="http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e">http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e</a></li> </ul>   | Upgrade preparation and implementation                            | <ul style="list-style-type: none"> <li>Preparing system to upgrade</li> <li>Installing software</li> <li>Performing upgrade tasks</li> </ul>                                    |
| <ul style="list-style-type: none"> <li>System Administration Guide at <a href="http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView">http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView</a></li> <li>Performance and Tuning Guide at <a href="http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView">http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView</a></li> </ul> | Planning/preparation prior to migration and testing after upgrade | <ul style="list-style-type: none"> <li>Planning system design for Adaptive Server 12.5</li> <li>Monitoring and tuning system for increased performance</li> </ul>               |

For additional information about Sybase database products, go to product manuals at <http://manuals.sybase.com/onlinebooks/group-as/asg1250e>

## Other Sources of Information

Use the Sybase Technical Library CD and the Technical Library Product Manuals Web site to learn more about your product:

- Technical Library CD contains product manuals and is included with your software. The DynaText browser (downloadable from Product Manuals at <http://www.sybase.com/detail/1,3693,1010661,00.html>) allows you to access technical information about your product in an easy-to-use format.

Refer to the *Technical Library Installation Guide* in your documentation package for instructions on installing and starting the Technical Library.

- Technical Library Product Manuals Web site is an HTML version of the Technical Library CD that you can access using a standard Web browser. In addition to product manuals, you will find links to the Technical Documents Web site (formerly known as Tech Info Library), the Solved Cases page, and Sybase/Powersoft newsgroups.

To access the Technical Library 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.

### ❖ For 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.
- 4 Select the Certification Report filter, specify a time frame, and click Go.
- 5 Click a Certification Report title to display the report.

### ❖ For the latest information on EBFs and Updates

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Select EBFs/Updates. Enter user name and password information, if prompted (for existing Web accounts) or create a new account (a free service).
- 3 Specify a time frame and click Go.
- 4 Select a product.

5 Click an EBF/Update title to display the report.

❖ **To create 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.

## Migration Web Page

For upgrade and migration-specific TechNotes and white papers, as well as links to other resources, see the ASE Migration Resources Web page at <http://sybase.com/support/techdocs/migration>.

## Classes

See the Sybase Education page for information about courses on Sybase products. For ASE 12.5 courses, see the Database Servers at <http://www.sybase.com/education/coursecatalog/databaseservers> page.

## Consulting Services

Sybase Consulting offers services to customers planning to migrate to Adaptive Server 12.5. Sybase Consulting uses the Sybase-developed and recommended methodology, SAFE/EM (Sybase Advanced Framework to Enable Effective Migration). The migration technology in this manual is based on SAFE/EM.

See the Consulting Web page at <http://www.sybase.com/solutions/services> or the ASE Migration Resources Web page at <http://sybase.com/support/techdocs/migration> to contact Sybase Consulting.

## **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.

# Analyze: Documenting Business Requirements

## In This Chapter

In this chapter you begin the first phase of migration planning, documenting your environment. This chapter helps you organize business information required for an effective migration plan. The sections are:

- Diagramming the System
- Operational Business Requirements
- Current Performance Metrics
- Additional Business Requirements

See Appendix A, “Worksheets for Your Current Environment” for worksheets like the ones used in the examples in this chapter.

## Diagramming the System

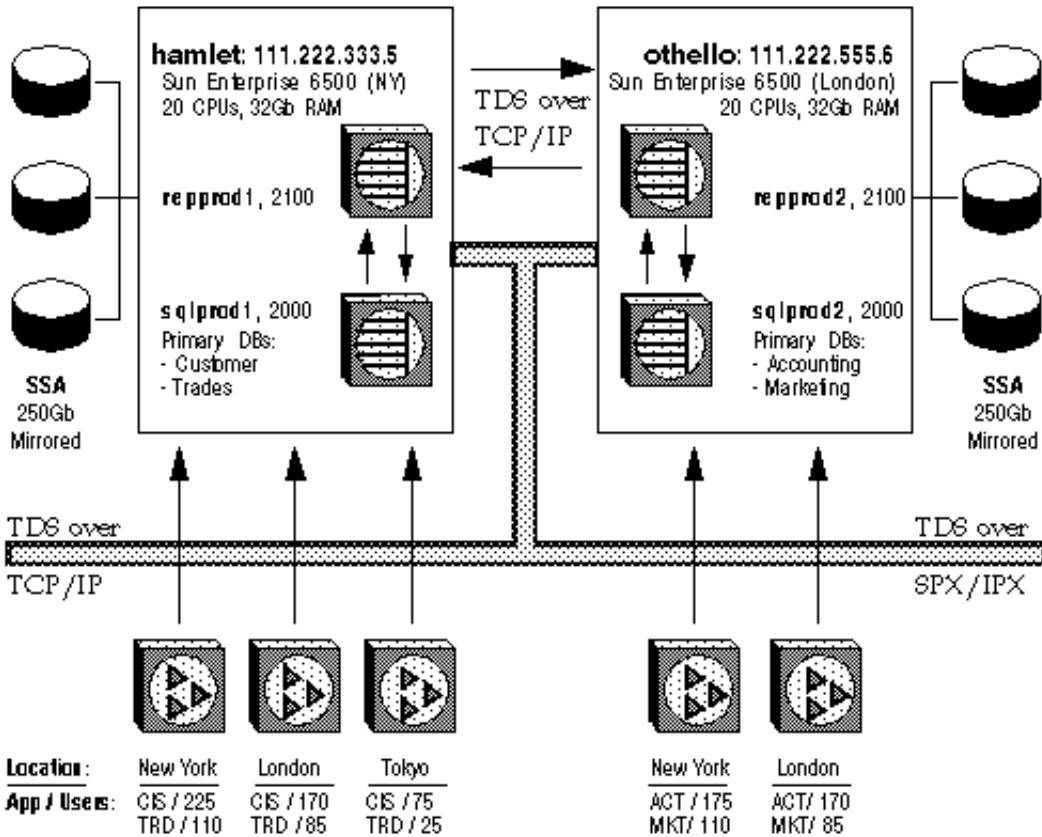
Create a diagram (or table or written description) illustrating the flow of information on your system. This will serve as a reference for the migration team. Include the following information:

- Servers, including file, print and application servers:
  - Machine name
  - IP address
  - Sybase name and aliases
- Clients
  - Applications
  - Number of users
- Network

- Protocols
- Gateways
- Routers, brouters, bridges

For example, you can create a diagram like the this high-level view of the Acme Brokerage, a firm with business in several cities and two main computing centers:

*Acme Brokerage*  
System Architecture





You can write a high-level business description in addition to or instead of a diagram.

## Operational Business Requirements

This section suggests ways to document your operational business requirements. You will use these baseline requirements to help you plan the migration and develop success criteria. This section covers the following issues:

- Availability Requirements
- Database Change Metrics
- Database Dump Details
- Maintenance Procedures
- Service Level Requirements
- Transaction Profile

## Availability Requirements

Record the times users need to access databases, as well as maximum acceptable downtime, as in the following example:

| <b>Database name</b> | <b>Operational hours</b>    | <b>Maximum downtime</b> | <b>General Comments</b> |
|----------------------|-----------------------------|-------------------------|-------------------------|
| TRD                  | 07:00 – 23:00<br>Mon to Fri | 5 min                   |                         |
| CIS                  | 07:00 – 23:00<br>Mon to Fri | 15 min                  |                         |
| ACT                  | 07:00 – 21:00<br>Mon to Sat | 5 min                   |                         |
| MKT                  | 07:00 – 20:00<br>Mon to Fri | 30 min                  |                         |

## Database Change Metrics

For all databases, record:

- Database size
- Transaction log growth
- Table rowcounts and daily change rate (number of inserts, deletes, and updates)

## Database Dump Details

Document your dump procedures, including times and devices, as in the following example:

| Database name | Frequency of database dumps | Dump device used       | Frequency of transaction log dumps | Dump device used | Comments |
|---------------|-----------------------------|------------------------|------------------------------------|------------------|----------|
| master        | every night                 | master_dumpdev         |                                    |                  |          |
| TRD           | every night                 | TRD_tape1<br>TRD_tape2 | Every 15 min.                      | TRD_tape2        |          |
| CIS           | every night                 | CIS_tape1              | Every 15 min                       | CIS_tape3        |          |
| ACT           | every night                 | ACT_tape1<br>ACT_tape2 | Every 15 min                       | ACT_tape3        |          |

## Maintenance Procedures

Document your schedule for running data consistency checking and performance monitoring tools. Use a table similar to this example:

| Database name | Frequency of dbcc checkdb and dbcc checktable | Frequency of dbcc checkalloc and dbcc tablealloc | Frequency of update statistics                | Frequency of monitoring utilization |
|---------------|---|--|---|-------------------------------------|
| master        | every night                                   |  |   |                                     |
| TRD           | every weekend                                 | different tables<br>“Round-robin” every night    | different tables<br>“Round-robin” every night | every hour                          |

| Database name | Frequency of dbcc checkdb and dbcc checktable | Frequency of dbcc checkalloc and dbcc tablealloc | Frequency of update statistics                | Frequency of monitoring utilization |
|---------------|---|--|---|-------------------------------------|
| CIS           | every weekend                                 | different tables<br>“Round-robin” every night    | different tables<br>“Round-robin” every night | every hour                          |
| ACT           | every weekend                                 | different tables<br>“Round-robin” every night    | different tables<br>“Round-robin” every night | every hour                          |

## Service Level Requirements

Document application details and service requirements, as in this example:

| Application name | Type of application  | App. langs    | Client machines    | No. of concurrent users | DBs accessed | Availability requirements (per day) | Performance (average response time) |
|------------------|----------------------|---------------|--------------------|-------------------------|--------------|-------------------------------------|-------------------------------------|
| Trades           | Heavy OLTP           | C             | UNIX work-stations | 220                     | TRD CIS      | <5 min downtime                     | <2 sec                              |
| Customer         | Light OLTP DSS       | Power-builder | PC                 | 470                     | CIS          | <10 min downtime                    | < 5 sec                             |
| Accounting       | Light OLTP DSS Batch | Power-builder | PC                 | 345                     | ACT CIS      | <5 min downtime                     | <5 sec                              |
| Marketing        | DSS                  | C             | PC                 | 195                     | MKT CIS      | <30 min downtime                    | <120 sec                            |

## Transaction Profile

Using statistics io, showplan, and both dbcc 302 and dbcc 310 to capture application processing details, record transaction profiles. You can plot these on a chart like the following. Save showplan and dbcc output for critical transactions. You will use this information as a baseline for post-upgrade testing, as described in Chapter 7, “Test: Ensuring Stability and Performance”.

Create a document similar to the following:

| App name | process (xact) name | Type of proces sing | Xact priority | Freq per user (per hour) | Source code          | Required average response time | Required max response time | Current avg/max response time |
|----------|---------------------|---------------------|---------------|--------------------------|----------------------|--------------------------------|----------------------------|-------------------------------|
| Trades   | addTrade            | Heavy OLTP          | P1            | 90                       | Stored Proc          | <2 sec                         | <5 sec                     | 1 sec avg 3 sec max           |
| Trades   | bustTrade           | Heavy OLTP          | P1            | 10                       | Stored Proc          | <5sec                          | <10 sec                    | 2 sec avg 8 sec Max           |
| Trades   | reconcileTrades     | Batch               | P1            | 1 per day                | Embedd ed SQL/CO BOL | <30 min                        | <60 min                    | 25 min avg 45 min max         |
| Trades   | listAccounts        | Light OLTP          | P1            | 180                      | Stored Proc          | <2 sec                         | <5 sec                     | 1 sec avg 2 sec max           |

---

**Note** If you are upgrading from ASE 12.0, you can save abstract query plans for your critical queries. For more information on abstract query plans, see the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView).

---

See the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView) for more information on gathering transaction statistics. Also see ASE Migration Resources Web page at <http://sybase.com/support/techdocs/migration> for TechNotes and white papers on query processing.

## Current Performance Metrics

Record as much of the following performance information as possible. Sybase monitors include ASE Monitor, Historical Monitor, and sp\_sysmon.

- CPU utilization:

Use operating system and Sybase monitors to measure the average and maximum CPU utilization (aggregate and per CPU on SMP servers) per “time window” (for example, online or batch) per server.

- Disk I/O:
  - Use operating system monitors to measure I/Os per second per disk and controller, and I/O queue lengths per “time window” per server.
  - Use Sybase monitors to measure total I/Os, reads, and writes per second per Sybase device per “time window” per server.
- Concurrency:
  - Use Sybase monitors to determine the average lock contention.
- Network I/O:
  - Use Sybase monitors to record average execution rates for critical stored procedures.
  - Use operating system monitors to measure the packets per second per network interface card per “time window” per server.
  - Use Sybase monitors to measure the TDS packets (“sent from” and “received by”) per “time window” per server.
- Memory:
  - Use operating system monitors to determine the paging/swapping rates per second per “time window” per server.
  - Use Sybase monitors to determine the data and procedure cache hit rates per “time window” per server.

See the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView) for information about using Sybase tools and stored procedures to monitor performance factors.

## Additional Business Requirements

Document any other important operational business requirements, such as:

- Priority list of applications to be migrated
- Constraints
  - Need to avoid year/quarter end processing?
  - How much downtime is ok?

- Should the upgrade take place over the weekend
- What staff is available: DBAs, system administrators?
- What resources are available: Replication Servers, machines, tools, funding?
- Application and data server dependencies
  - Is more than one application using the same ASE?
  - Can all applications on a server be migrated?
- Are there any vendor issues? Are your third-party applications certified to run against ASE 12.5?

# Analyze: Documenting Your Environment

## In This Chapter

This chapter provides guidelines for documenting system hardware and software in the Adaptive Server production environment. This information is used to identify resource issues during the planning phase of the migration.

See Appendix A, “Worksheets for Your Current Environment” for worksheets like the ones used in the examples in this chapter.

This chapter contains the following sections:

- Hardware Configuration
- Physical Memory Utilization
- Software Configuration
- Sybase Configuration
- Adaptive Server Objects

## Hardware Configuration

Document your hardware environment as described in the following sections:

- General Server Hardware
- CPU Resources per Machine
- Disk Configuration
- Network Configuration
- Tape Configuration

## General Server Hardware

List the following hardware information for every server machine:

- Make and model
- Your customer ID with the vendor
- Technical support information
  - Telephone number
  - Support hours
  - Name of your account manager and his or her telephone or pager number
- Vendor's Web page

## CPU Resources per Machine

List the following CPU information for each server machine:

- Total number of processors and their speed
- Number of processors available to ASE
- Other CPU-intensive processes sharing those processors
- List of processes/threads bound to specific CPUs
- List of processes/threads run at high priority

## Disk Configuration

Use worksheets like those in the following examples to gather the following disk I/O information:

- Controller map
- Disk layout map
- Disk partition map
- Logical volume map

---

| <b>Controller Number</b> | <b>Make and Model</b> | <b>Firmware Revision</b> | <b>Months in Service</b> | <b>Transfer Rate (Kb/sec)</b> |
|--------------------------|-----------------------|--------------------------|--------------------------|-------------------------------|
| 0                        | Sun Sparc 5           | 1.50                     | 9                        | 7500                          |

---



| Controller Number | Make and Model | Firmware Revision | Months in Service | Transfer Rate (Kb/sec) |
|-------------------|----------------|-------------------|-------------------|------------------------|
| 1                 | Sun Sparc 20   | 1.00              | 6                 | 10000                  |

| Physical Device name | Make and Model   | Firmware Revision | Months in Service | Controller Number | Capacity (Mb) | Throughput (I/Os per sec) | Transfer Rate (Kb/sec) |
|----------------------|------------------|-------------------|-------------------|-------------------|---------------|---------------------------|------------------------|
| c0t0d0               | Seagate ST43401N | 2.15              | 9                 | 0                 | 2900          | 80                        | 1500                   |
| c0t0d1               | Seagate ST43401N | 2.15              | 12                | 0                 | 2900          | 80                        | 1000                   |
| c0t0d2               | Seagate ST43401N | 2.15              | 24                | 0                 | 2900          | 80                        | 1600                   |
| c0t0d3               | Seagate ST43401N | 2.00              | 16                | 0                 | 2900          | 80                        | 1200                   |

| Logical Volume Name | Member Disk Partitions                       | Used by (Sybase, UFS) | Sybase Device Name | Mirrored Logical Device | Capacity (Mb) | Stripe Width (Mb) |
|---------------------|--|-----------------------|--------------------|-------------------------|---------------|-------------------|
| lv dev1             | c0t0d0s3<br>c0t0d0s4<br>c0t0d1s3<br>c0t0d1s4 | sybase                | TRD Log            | lv dev1 mirror          | 500           | 4                 |
| lv dev2             | c0t0d0s3<br>c0t0d0s4<br>c0t0d1s3<br>c0t0d1s4 | sybase                | CIS Log            | lv dev2 mirror          | 500           | 4                 |

| Physical Device Name | Partition Number | Used by (Sybase, UFS) | Device Name | OS Mirror Device Name | Capacity (mb) | Cylinder Range |
|----------------------|------------------|-----------------------|-------------|-----------------------|---------------|----------------|
| c0t0d0               | s0               | disk label            |             |                       | 2             | 0 – 1          |
|                      | s2               | backup                |             |                       |               |                |
|                      | s3               | swap                  | swap        |                       | 998           | 2 – 501        |
|                      | s4               | sybase                | TRD Log     | c1t0s4                | 500           | 502 – 752      |
|                      | s5               | sybase                | CIS Log     | c1t0s5                | 500           | 753 – 1003     |
|                      | s6               | ufs                   | /usr        |                       | 900           | 1004 – 2733    |

| Physical Device Name | Partition Number | Used by (Sybase, UFS) | Device Name | OS Mirror Device Name | Capacity (mb) | Cylinder Range |
|----------------------|------------------|-----------------------|-------------|-----------------------|---------------|----------------|
| c0t0d2               | s0               | disk label            |             |                       | 2             | 0 – 1          |
|                      | s2               | backup                |             |                       |               |                |
|                      | s3               | swap                  | swap        |                       | 2900          | 2 – 2733       |

## Network Configuration

Use a worksheet like the one in the following example to show network interface card information for server and client machines:

| Physical Device Name | Make and Model | Firmware Revision | Months in Service | Protocols supported | Network Address | Transfer Rate (Kb/sec) |
|----------------------|----------------|-------------------|-------------------|---------------------|-----------------|------------------------|
| c0t0d0               | Sun Sparc 5    | 1.5               | 9                 | TCP/IP<br>SPX/IPX   | 121.222.233.11  | 7500                   |
| c0t0d1               | Sun Sparc 20   | 1.00              | 12                | TCP/IP<br>SPX/IPX   | 121.222.555.33  | 7500                   |

## Tape Configuration

Use a worksheet like the one in the following example to record tape or other storage media configuration:

| Physical Device Name | Make and Model | Firmware Revision | Months in Service | Controller Number | Capacity (Mb_) | Transfer Rate (Kb/sec) |
|----------------------|----------------|-------------------|-------------------|-------------------|----------------|------------------------|
| /dev/rmt/0           | Sun Sparc 10   | 2.15              | 9                 | 2                 | 2000           | 500                    |
| /dev/rmt/1           | Sun Sparc 20   | 1.00              | 12                | 2                 | 2900           | 500                    |

## Physical Memory Utilization

List all the major processes running on a server machine and use the formulas given here to calculate their memory requirements. Add the individual sizes together for total memory requirements.

Use this table as a guide:

| <b>Name</b>                   | <b>Runtime Memory Usage Calculation</b>   |
|-------------------------------|---|
| Operating System              | OS-specific   |
| Adaptive Server               | Add together these ASE parameters:<br>total memory + additional netmem + extent i/o buffers |
| Backup Server                 | Add:<br>size of the binary + (110Kb * number of stripes)                                    |
| sybmultbuf                    | Size of the binary only (attaches to BS shared memory)                                      |
| Replication Server            | See “memory limit” parameter  |
| Monitor Server                | Size of the binary only   |
| Gateways {List:}              | Specific to the product - see product documentation   |
| Other Applications<br>{List:} | Specific to the application   |
| Total Memory Required         |   |
| Total Memory Installed        |   |

## Software Configuration

Document your software environment as described in the following sections:

- Operating System
- Applications

### Operating System

List the following operating system information:

- Operating system name
- Release level
- Patch level
- Kernel configuration parameters
- Swap size
- OS-specific software installed

- High availability software installed

You may need to contact your operating system vendor to get system upgrades, recent patches, or help with problems, so make a note of the following operating system technical support information:

- Telephone number
- Support hours
- Name of your technical account manager and his or her telephone or pager number
- Vendor's Web page

## Applications

Make a list of applications to be migrated to ASE 12.5. For each application, note:

- Information about data and usage
  - Distributed data
  - Warehoused or used in transaction processing? If used in transaction processing, it needs to be accurate and in the right format. Warehoused data is more tolerant of format and slight computational differences, such as datatype conversion differences.
- Location of application source files
- Types and number of modules containing SQL code to be evaluated for needed modifications (triggers, stored procedures)

## Sybase Configuration

Document your Sybase configuration as described in the following sections:

- General Information
- Database Devices
- Databases and Segments
- Dump Devices

## General Information

Record the following Sybase information:

- Adaptive Servers and their SYBASE home directories
- Components and release levels (including EBFs)
- Names and location of scripts to rebuild database environment
- All server configuration values; these can be listed by `sp_configure` or found in the `.cfg` file.
- Server runtime memory map: Run `dbcc memusage` during non-peak time or in single user mode. This command tells you how memory is configured within ASE. See the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView) for more information about ways to capture server memory statistics. You can record the output of `dbcc memusage` in a table like this one:

|                         | <b>Mb</b> | <b>2K<br/>Blocks</b> | <b>Bytes</b> |
|-------------------------|-----------|----------------------|--------------|
| Configured Memory       | 10.0000   | 5120                 | 10485760     |
| Code size:              | 2.7075    | 1387                 | 2839060      |
| Kernel Structures:      | 1.5924    | 816                  | 1669755      |
| Server Structures:      | 1.8529    | 949                  | 1942928      |
| Page Cache:             | 3.0496    | 1562                 | 3197756      |
| Proc Buffers:           | 0.0270    | 14                   | 28348        |
| Proc Headers:           | 0.6191    | 317                  | 649216       |
| Number of page buffers: |           | 1489                 |              |
| Number of proc buffers: |           | 372                  |              |

## Database Devices

Record database device information, as in the following example:

| <b>Database Device<br/>Name</b> | <b>Physical Device name</b> | <b>Mirrored<br/>Device<br/>Name</b> | <b>Virtual Device<br/>Number</b> | <b>Size (Mb)</b> |
|---------------------------------|-----------------------------|-------------------------------------|----------------------------------|------------------|
| TRD_dev1                        | /dev/rdisk/c0t0d0s3         |                                     | 2                                | 10020            |
| TRD_dev2                        | /dev/rdisk/c0t1d0s3         |                                     | 3                                | 5020             |
| TRD_log                         | /dev/rdisk/c0t1d0s4         |                                     | 4                                | 1020             |
| CIS_dev1                        | /dev/rdisk/c0t1d1s3         |                                     | 5                                | 4020             |

| Database Device Name | Physical Device name | Mirrored Device Name | Virtual Device Number | Size (Mb) |
|----------------------|----------------------|----------------------|-----------------------|-----------|
| CIS_log              | /dev/rdisk/c0t1d1s4  |                      | 6                     | 420       |

## Databases and Segments

List of all segments and the objects on them. Use a worksheet like that in the following example:

| Database Name | DB Options Set       | Size (Mb) | Segments Names             | Device Name | Size (Mb) |
|---------------|----------------------|-----------|----------------------------|-------------|-----------|
| master        | none                 | 700       | default.system.log         | master      | 3         |
| master        | none                 | 500       | default.system.log         | master      | 2         |
| master        | none                 | 300       | default.system.log         | master      | 1         |
| model         | none                 | 200       | default.system.log         | master      | 2         |
| tempdb        | select into/bulkcopy | 200       | default.system.log         | master      | 2         |
| TRD           | none                 | 10000     | system, trd_seg1           | TRD_dev1    | 10020     |
| TRD           | none                 | 5000      | system, trd_seg1, trd_seg2 | TRD_dev2    | 5020      |
| TRD           | none                 | 1000      | log                        | TRD_log     | 1020      |
| CIS           | none                 | 4000      | system, cis_seg1           | CIS_dev1    | 4020      |
| CIS           | none                 | 400       | log                        | CIS_log     | 420       |

## Dump Devices

Record dump device information as in the following example:

| Database Device Name | Physical Device name | Media type | Capacity (Mb) |
|----------------------|----------------------|------------|---------------|
| Tape dev1            | /dev/rmt/0m          | 4mm        | 2000          |
| Tape dev2            | /dev/rmt/1m          | 4mm        | 2000          |
| Tape dev3            | /dev/rmt/2m          | 4mm        | 2000          |
| Tape dev4            | /dev/rmt/3m          | 4mm        | 2000          |
| Tape dev5            | /dev/rmt/4m          | 4mm        | 2000          |

## Adaptive Server Objects

Document the objects in your current Adaptive Server as described in the following sections:

- Gather Scripts to Create Objects
- If You Don't Have Scripts

### Gather Scripts to Create Objects

Locate or create the scripts necessary to recreate:

- Server level objects
  - Database devices
  - Configurations
  - Logins and security
- Database level objects, including:
  - Defaults, rules, and user datatypes
  - User databases
  - Users, groups, and aliases
  - Tables, views, and stored procedures
  - Other database objects such as triggers and indexes

You may also want need to extract and load data with bcp. These scripts can be used to help set up your test environment as well as the building a new production system and may be needed if you plan to maintain two server systems at different release levels.

### If You Don't Have Scripts

There are several ways you can reproduce scripts or access the information needed to reproduce configuration and objects.

## Query Sybase System Tables

The following system tables contain object information that you can use to create installation scripts:

- sysdatabases
- sysdevices
- sysusages
- sysobjects
- sysusers
- syssservers
- syslogins
- sysremotelogins

See the System Administration Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView) for more details on system tables and objects. See also Segment Remapping with load database When Moving a Database at <http://my.sybase.com/detail?id=1324> in Technical Information Library for information on using system tables to reconstruct databases.

## Use System Stored Procedures

For information about current ASE configuration, use `sp_configure` with no argument. It will list all configuration parameters and their values.

For the SQL commands in a stored procedure, use `sp_helptext`.

See the Reference Manual at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/refman/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/refman/@Generic__BookView) for other system stored procedures such as `sp_helpdevice`, `sp_help`, `sp_helpdb`, `sp_helpsegment` and `sp_helpindex` that can provide information about objects on your server.

## Use Sybase Tools

You can reverse engineer server objects using Sybase DBA tools such as Sybase Central, or Powersoft PowerDesigner, or use third-party tools.



# Plan: Writing a Plan and Getting Ready to Migrate

## In this Chapter

Now that you have collected data about your current system, you can write a migration plan. This chapter discusses migration methods and planning. In addition to choosing a migration method, you may need to bring your system resources to the level required by ASE 12.5 and make changes needed in applications and system administration procedures.

This chapter contains the following sections:

- Review Documentation
- Review the Upgrade Process
- Determine Migration Approach
- Write a Migration Plan
- Build the Adaptive Server Environment

## Review Documentation

Before planning and carrying out the migration, become familiar with Adaptive Server 12.5 migration issues. Good starting points include:

- Information available online on the ASE Migration Resources Web page at <http://sybase.com/support/techdocs/migration>, including TechNotes, white papers, and compatibility information.
- What's New in Adaptive Server Enterprise at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic__BookView)

- The Release Bulletin at [http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic\\_\\_CollectionView;pt=asp1250e](http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e) for your platform
- The installation guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic\\_\\_CollectionView;pt=asp1250e](http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e) for your platform

For other documentation and resources, see Chapter 1, “Introduction and Guide to Resources.”

## Review the Upgrade Process

Changes in the upgrade process were introduced in Releases 11.9.2 and 12.0. These changes include:

- The Studio Installer installation utility
- Component licensing through Sybase Software Asset Management (SySAM)
- New directory structure
- New environment variables

See the section on *The Upgrade Process* in the white paper Technical Issues in ASE 11.9.x and 12.0 Upgrade at <http://www.sybase.com/detail/1,3693,1011767,00.html> for more information about changes in the upgrade process as well as recommendations for upgrading. For information on SySAM, see the white paper License Manager - Sybase Software Asset Management at <http://my.sybase.com/detail?id=1009350>. If you intend to upgrade to the latest bug-fix release, see the following TechNotes:

- Loading Instructions for Adaptive Server Enterprise 12.5 EBFs at <http://www.sybase.com/detail?id=1013599>
- Post-installation Scripts and Tasks for Adaptive Server Enterprise 12.5 EBFs at <http://www.sybase.com/detail?id=1013610>

## Special Cases: Migrating to 64-bit Operating System or Larger Page Size

As of release 11.9.3, Adaptive Server runs on 64-bit operating systems. As of 12.5, ASE allows you to create a server that uses a logical page size larger than 2K in order to store more data per page. These changes are described in What's New in Adaptive Server Enterprise at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic__BookView) and the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView).

If you wish to take advantage of these changes, you need to select the appropriate migration method. These are the options currently available:

| Type of migration                                | Methods/tools available   |
|--|---|
| Pre-12.5 32-bit server to 12.5 64-bit server     | <p>You can use the <code>sqlupgrade</code> utility.</p> <hr/> <p><b>Note</b> You can use the <code>sqlupgrade</code> utility only when you are changing release levels. You cannot use it just to move from a 32-bit to a 64-bit system.</p> <hr/>  |
| 12.5 32-bit system to 12.5 64-bit system         | <p>Create a new 12.5 installation and use <code>bcp</code> or <code>dump</code> and load to move data into the new server. You can also manually replace the binary to change to a 64-bit system. See the installation guide at <a href="http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e">http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e</a> for instructions.</p>   |
| Pre-12.5 2K page server to 4K, 8K, or 16K server | <p>As of the release of this manual, an upgrade to a larger page size is not supported. You need to rebuild the system using the instructions for 32-bit to 64-bit migration in the installation guide at <a href="http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e">http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e</a>. Watch the ASE Migration Resources Web page at <a href="http://sybase.com/support/techdocs/migration">http://sybase.com/support/techdocs/migration</a> for developments on tools to aid data migration to larger page size.</p> |

## Determine Migration Approach

The best migration strategy for you depends on such factors as the cost of the effort, the type of business you do, the size of your databases, and available resources.

The following table highlights the advantages and disadvantages of each migration approach:

| Approach                    | Advantages  | Disadvantages   | When used   |
|-----------------------------|---|---|---|
| Parallel With Replication   | Easy fallback to earlier release. You do not need to rebuild previous release databases.<br>Minimal system down time. | Can be complex in OLTP environments.<br>Replication Server must be set up, requiring extra hardware and software.                                     | This approach is best for large 7x24 production databases, maintaining high availability, when: <ul style="list-style-type: none"><li>• Rebuilding a release database can take too long.</li><li>• The system may have a large number of transactions and complex Transact-SQL queries with subqueries.</li></ul> |
| Cutover Without Replication | Can be executed with minimal resource demands.  | Highest risk. Requires down time for critical migration tasks.<br>Recovery can be time consuming in a production environment.                         | This approach is good for resource-constrained environments. It is appropriate for large organizations only if you are able to schedule sufficient downtime, for instance, a long weekend.  |
| Phased Cutover              | Low risk with low development overhead.<br>Especially conducive to testing.   | May require additional resources—either more memory or a second system.<br>Requires tighter coordination with application groups and database owners. | If neither of the other two approaches seems appropriate, you can use a phased cutover.   |

For more information on these approaches, see the following sections:

- Parallel With Replication
- Cutover Without Replication

- Phased Cutover

---

**Note** This migration guide does not cover other parallel migration approaches; such as running two systems in parallel where you have to maintain both the systems simultaneously, or transaction duplication where you use one front-end to drive two parallel back-ends. These system operational approaches include factors too site-specific to detail effectively in this guide.

Whenever possible, upgrade *test* and *development* databases to release 12.5 first. Upgrade the production system after testing. See Chapter 7, “Test: Ensuring Stability and Performance” for more information about testing.

---

## Parallel With Replication

The issues for the parallel-with-replication approach are described in the following sections:

### Method

The following general method can be used to perform a parallel upgrade with replication:

- 1 Install a new copy of ASE 12.5
- 2 Copy in or load pre-release 12.5 databases
- 3 Use Replication Server to maintain both sets of databases. The 12.5 system becomes the primary server, and you maintain the pre-12.5 system as a warm standby.

### Fallback

Plan for all users to reconnect to the earlier version server after you take ASE 12.5 offline. Make TCP/IP address and port changes where appropriate.

Test the fallback process as part of the application test suite. This suite should do both of the following:

- Insert data into ASE 12.5. The data must be replicated and available in the earlier server.
- Execute the fallback script.

Consider making daily bcp dumps of the databases. To fall back, load the dumps into the earlier server. Keep in mind:

- You may need to modify the databases to support incremental bcp dumps.

- The earlier server cannot read release 12.5 backup files. You need to create bcp or other scripts to move tables back to pre-release 12.x.
- Do not apply schema enhancements.

For information about scheduling backups of user databases, see the System Administration Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView).

Additional tips:

- Upgrade Replication Server first.
- Be sure the applications are using the correct server. For details on the interfaces file and the \$DSQUERY environment variable, see the configuration guide for your platform.
- Remember to include client fallback time in the calculation.
- For 24x7 sites, load time delays can impact synchronization. Consider replicating to ASE 12.5 and then switching servers.

Application test suite

Since the parallel with replication approach is best for high availability applications, it is imperative that the test suite address both update correctness and performance acceptability. Execute this suite *before* switching users to the new system.

See Chapter 7, “Test: Ensuring Stability and Performance” for more information on testing.

**Note** After successful validation, consider having users enter production queries with the production toolset. A good time to do so is after hours or during production lulls.

Bridging

There should not be any user impact during migration. The more stringent the validation test is, the less likely you are to have bridging issues.

To ensure correct updates and acceptable performance, test the replicated environment.

See the Replication Server product manuals at [http://www.sybase.com/detail\\_list/1,3691,2721,00.html](http://www.sybase.com/detail_list/1,3691,2721,00.html) for more information.

Environment

The environment used for ASE 12.5 needs to be more powerful to handle the query and replication loads. See the Replication Server configuration guide for your platforms by going to the Replication Server product manuals at [http://www.sybase.com/detail\\_list/1,3691,2721,00.html](http://www.sybase.com/detail_list/1,3691,2721,00.html).

Be sure to account for any increased release 12.5 memory requirements that apply to your configuration. For more information, see:

- The installation guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic\\_\\_CollectionView;pt=asp1250e](http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e) for your platform. It gives basic RAM requirements.
- Chapter 6, “Implement: Making Database Administration Changes”.
- Details on configuring memory and data caches in the System Administration Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView)
- Information on how to configure memory for performance in the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView).

**Note** For a production system, execute the performance suite during off hours.

Scheduling

Developing and running the replication facility, validation and performance suite, and fallback script requires significant effort. If your environment already uses replication, this effort will be notably easier.

For a *development* system, you may want to add a short period to the development schedule for release 12.5 issues.

For a *production* system, be prepared to postpone or back off if needed.

## Cutover Without Replication

The issues for the cutover without replication approach are described in the following sections:

Method

Upgrade all databases to release 12.5 at the same time. A cutover without replication is common in small organizations for development or production servers. It may be used in larger organizations when you are able to schedule sufficient downtime.

Fallback

Base fallback on the amount of time needed to restore earlier databases. For example, if users need the system Monday at 8 a.m. and the restoration takes 8 hours, the validation test must pass by Sunday at midnight.

---

**Note** Remember to include client fallback time in the calculation.

---

|  |   |
|--|---|
|  | <p>You can use dump database or bcp out before an upgrade to prepare for fallback.</p> <p>Plan a way to capture transactions after cutover to be used in case of fallback. If you go into production and then need to back off, you will have to restore all the transactions that occurred after the last dump/load.</p>   |
| Application test suite   | <p>For a development system, simple validation may be adequate. However, for a <i>production</i> system, the test suite must address both update correctness and performance acceptability.</p> <p>See Chapter 7, “Test: Ensuring Stability and Performance” for more information on testing.</p>   |
| <hr/> <b>Note</b> You might want to validate over a 3-day weekend if possible. <hr/> |   |
| Fallback after cutover   | <p>Consider making daily bcp dumps of the databases. You can then load the dumps into the earlier server to fall back. Keep in mind:</p> <ul style="list-style-type: none"><li>• You may need to modify the databases to support incremental bcp dumps.</li><li>• Earlier servers cannot read release 12.5 backup files. You need to create bcp or other scripts to move tables back to pre-release 12.5.</li><li>• Do not apply 12.5 schema enhancements.</li></ul> <p>For information about scheduling backups of user databases, see the System Administration Guide at <a href="http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView">http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView</a>.</p>   |
| Bridging   | <p>There should not be any user impact during migration. The more stringent the validation test is, the less likely you will have bridging issues.</p>  |
| Environment  | <p>Be sure to account for any increased release 12.5 memory requirements that apply to your configuration. For more information, see:</p> <ul style="list-style-type: none"><li>• The installation guide at <a href="http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e">http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e</a> for your platform. It gives basic RAM requirements.</li><li>• Chapter 6, “Implement: Making Database Administration Changes”.</li><li>• Details on configuring memory and data caches in the System Administration Guide at <a href="http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView">http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView</a></li></ul> |



- Information on how to configure memory for performance in the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView).

---

**Note** For a production system, execute the performance suite during off hours.

---

Scheduling

For a *development* system, you may want to add a short period to the development schedule for release 12.5 issues.

For a *production* system, be prepared to postpone or back off if needed.

## Phased Cutover

The issues for a phased cutover are described in the following sections:

Method

Change only one application and database to release 12.5 at a time.

Fallback

Consider making daily bcp dumps of the databases. To fall back, you can then load the dumps into the earlier server. Keep in mind:

- You may need to modify the databases to support incremental bcp dumps.
- The earlier server cannot read release 12.5 backup files. You need to create bcp or other scripts to move tables back to the earlier server.
- Do not use release 12.5 schema enhancements until the conversion succeeds.

For information about scheduling backups of user databases, see the System Administration Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView).

Application test suite

Ensure that the application test suite addresses both update correctness and performance acceptability. Also ensure that you do the following:

- Maintain the directories/libraries for both releases.
- Make sure the applications are using the correct server.
- After successful validation, consider having users enter production queries with the production toolset. A good time to do so is after hours or during production lulls.

See Chapter 7, “Test: Ensuring Stability and Performance” for more information on testing.

### Bridging

There should not be any user impact during migration. The more stringent the validation test is, the less likely you will have bridging issues.

The earlier server cannot read release 12.5 backup files. You need to create bcp or other scripts to move tables back to pre-release 12.5.

---

**Note** Do not use release 12.5 schema enhancements until the conversion succeeds.

---

### Environment

Be sure to account for any increased release 12.5 memory requirements that apply to your configuration. For more information, see:

- The installation guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic\\_\\_CollectionView;pt=asp1250e](http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e) for your platform. It gives basic RAM requirements.
- Chapter 6, “Implement: Making Database Administration Changes”.
- Details on configuring memory and data caches in the System Administration Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/sag/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/sag/@Generic__BookView)
- Information on how to configure memory for performance in the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/perf/@Generic__BookView).

Here are some additional tips:

- Execute performance measurements on a system with similar capabilities.
- For a production system, execute the performance suite during off hours.

### Scheduling

For a *development* system, you may want to add a short period to the development schedule for release 12.x issues.

For a *production* system, be prepared to postpone or back off if needed.

Be sure to notify users when you will be taking applications or databases offline.

## Write a Migration Plan

Produce a project plan which documents:

- Migration strategy—Which method is most appropriate for your site.
- Fallback—What to do in case the migration fails. The plan you evolve will be site-specific, but some general issues are discussed in Chapter 7, “Test: Ensuring Stability and Performance”.
- Application test suite—What validation and performance testing to perform for acceptance. See Chapter 6, “Implement: Making Database Administration Changes” for guidance.
- Bridging—Ways to minimize the impact to users during the migration. Refer to the business requirements you gathered in Chapter 3, “Analyze: Documenting Your Environment”.
- Environment—Additional resources and changes to the environment needed based on the information you gathered in Chapter 3, “Analyze: Documenting Your Environment”.
- Scheduling—How much time the migration will take based on the level of complexity and business needs. Refer to the business requirements you gathered in Chapter 3, “Analyze: Documenting Your Environment”.

Producing the following documentation may also be useful:

- A work breakdown that lists tasks chronologically and assigns them to specific roles like the one in Appendix B, “Sample Migration Task Lists”.
- Specification for application changes. The details of needed application changes are discussed in Chapter 5, “Implement: Making Required Application Changes”.

## Build the Adaptive Server Environment

After deciding the best migration approach for your system, begin creating or updating the environment for Adaptive Server 12.5. Go to the following sections and chapters for information about preparing your environment:

- Update Hardware Resources
- Verify Operating System Version and EBF Level
- Review Adaptive Server Interoperability with Other Sybase Products

- Update Applications and System Administration Procedures
- Create Migration Scripts
- Create a test environment. See Chapter 7, “Test: Ensuring Stability and Performance” for more information on the test environment

## **Update Hardware Resources**

Evaluate hardware resource needs according to your migration approach. For example, if you plan to use the parallel-with-replication method, you may need additional disk space to install a secondary system, a Replication Server for the high availability environment, a disk farm for backout strategies, and more memory.

The Release Bulletin at [http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic\\_\\_CollectionView;pt=asp1250e](http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e) for your platform contains information about any hardware requirements.

See the installation guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic\\_\\_CollectionView;pt=asp1250e](http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e) for your platform for basic RAM requirements. See Chapter 6, “Implement: Making Database Administration Changes” for more information on memory requirements.

See the System Administration Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/sag/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/sag/@Generic__BookView) or information on memory and backups.

## **Verify Operating System Version and EBF Level**

Ensure that the operating system is at the proper version and level to run Adaptive Server 12.5. Be sure that you have installed the most recent operating system patches to ensure that you have the latest bug fixes.

For release 12.5 operating system requirements, see the installation guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic\\_\\_CollectionView;pt=asp1250e](http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e) for your platform. See the ASE Migration Resources Web page at <http://sybase.com/support/techdocs/migration> for updates about recommended Sybase EBFs (bug fixes).

**Note** If you need to perform an operating system upgrade, do so before migrating. Test the new system to be sure it's working properly to avoid introducing unrelated errors into the migration process.

## Review Adaptive Server Interoperability with Other Sybase Products

To ensure that the versions of other Sybase products in use at your site are compatible with Adaptive Server 12.5, see Technical Documents at <http://sybase.com/support/techdocs> or the ASE Migration Resources Web page at <http://sybase.com/support/techdocs/migration> for recent interoperability and compatibility information.

## Update Applications and System Administration Procedures

The following chapters in this manual describe ASE 12.5 upgrade issues:

- Chapter 5, “Implement: Making Required Application Changes”
- Chapter 6, “Implement: Making Database Administration Changes”

Migrating to Adaptive Server 12.5 requires you to review your current applications and system administration procedures for any changes that could cause system problems or unexpected processing results.

This is the most time-consuming part of migration preparation; choose a method to do this that suits your needs and resources. Possible choices include:

- Writing your own scripts to examine and change applications
- Running the reserved word check from to 12.5 and using some of the tools available such as the stored procedures `sp_checkreswords` and `sp_procqmode`. See the Reference Manual at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/refman/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/refman/@Generic__BookView) for information about these stored procedures.

## **Create Migration Scripts**

Using the scripts you located, wrote, or reverse engineered in Chapter 2, “Analyze: Documenting Business Requirements”, write or edit the scripts that will create your 12.5 Adaptive Server installation. You need:

- Server-level migration scripts to create the new Adaptive Server environment, including database devices, configuration, logins, and security
- Database-level migration scripts to create Adaptive Server user databases and database objects such as tables, views, indexes, triggers, groups, users, and permissions.

# Implement: Making Required Application Changes

## In This Chapter

This chapter and Chapter 6, “Implement: Making Database Administration Changes” divide technical issues into those relevant to application developers and those relevant to database administrators. However, many issues are not exclusive to either of these roles and you should check both chapters whatever your role. This chapter covers those issues that may affect the execution of applications or that might require coding changes.

This chapter does not attempt to cover new features or changes that are not likely to cause surprises. For a comprehensive listing of changes and new features, see *What’s New in Adaptive Server Enterprise* at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic__BookView).

---

**Note** Changing applications and system administration procedures is the most time-consuming part of migration preparation and this guide does not tell you how to make these changes. You must choose a method to do this that suits your needs and resources. For instance, you may wish to develop your own plans and scripts, or you may prefer to contact Sybase Consulting for help. The ASE Migration Resources Web page at <http://sybase.com/support/techdocs/migration> contains migration TechNotes and white papers, and Sybase Technical Support can help with technical problems such as bugs and error conditions.

---

Begin reading at the earliest section that includes your current ASE version:

- If Your Version is 11.5
- If Your Version is 11.5 or 11.9.x
- If Your Version is 11.5, 11.9.x, or 12.0

## If Your Version is 11.5

This sections covers the following topics:

- Changes to Locking
- Changes to Error Checking
- 11.9.2 Query Processing Changes

### Changes to Locking

Release 11.9.2 introduced several changes to locking, including support for row-level and page-level locking. Some of these changes require you to make application or configuration changes to avoid performance issues.

#### Data-only Locking

Sybase now supports new locking methods called “data-only” locking. In previous releases, all locking happened at the page level. When you upgrade, by default, your server performs page-level locking. However, you may change the locking scheme on your existing tables with the `alter table` command.

The new locking schemes are:

- Datapages locking
- Datarows locking, also known as row-level locking

If you intend to use the data-only locking feature, you need to be aware of the following system and application issues:

- The number of locks required for your applications may increase, and you may need to change the *number of locks* configuration parameter.
- On data-only-locked tables with clustered indexes, the default behavior of disallowing duplicate rows and the `ignore_dup_row` option are not enforced during inserts and updates. This changes the behavior of commands that insert duplicate rows and may change the behavior of `create clustered index` if the table contains duplicate rows.
- On data-only-locked tables, the `sorted_data` option to create index can be used only immediately following a bulk copy operation that copies into an existing table. Use of the `sorted_data` option is prohibited once additional page allocation operations have been made for the table.



- Bulk copy into data-only-locked tables requires the version of bcp and the bulk copy libraries shipped with Adaptive Server version 11.9.2. Older versions of bcp and the bulk-copy libraries can still be used to copy into allpages-locked tables. See the Release Bulletin for the required version number.
- When using parallel sort for data-only-locked tables, the number of worker processes must be configured to equal or exceed the number of partitions, even for empty tables. The database option `select into/bulkcopy/pllsort` must also be enabled.

Clustered indexes on data-only-locked tables may not return rows in clustered key order if there is no order by clause. Bulk copy is not guaranteed to copy out a table in clustered key order.

Queries on unpartitioned allpages-locked tables with clustered indexes return rows in the clustered key order if the query does not include a sort (that is, if there are no clauses such as `order by` or `distinct`). Bulk copy also copies the rows out of allpages-locked tables in clustered key order for both partitioned and unpartitioned tables.

## Changes to Table-Level Locking

Performing an update or delete via a table scan does not acquire a table-level exclusive lock as it did in earlier versions. In version 11.9.2, the scan acquires an exclusive intent table lock when the transaction starts. When the scan locates a row that must be updated, it acquires an exclusive page or row lock, depending on the locking scheme. The only exception is an unindexed update or delete to a data-only-locked table at transaction isolation level 3; these updates acquire an exclusive table lock. In previous versions, any update or delete that does not use an index acquires a table-level exclusive lock at the start of the transaction.

In most cases, performing the scan without using a table level lock increases concurrency. However, some applications may experience one of these problems:

- Applications that did not deadlock in earlier versions may experience deadlocks in 11.9.2.
- Tasks that update a large number of rows may use a large number of locks if they cannot perform lock promotion due to conflicting locks.

## Workarounds

If deadlocks are creating problems, consider converting the table to use datarows locking.

If deadlocks or excessive numbers of locks are creating problems, consider using the lock table command to acquire an exclusive table lock before you begin the transaction.

## Changes to Error Checking

As of version 11.5.1, ASE performs more rigorous checking of permissions-related access to database objects. This change was not documented in 11.5 documentation. The changes may affect the behavior of applications that test for specific error codes and those that perform auditing of failed permissions. The types of changes include the following:

- Some permission checks that used to be performed at compile time are now delayed until execution time. This closes a loophole that allowed users to execute stored procedures if copies of the procedure remained in cache after permissions for the user were revoked. It can affect the type of error, and the error number, returned by some command batches.
- The ordering of checks for existence, checks for permissions, integrity checks, and the auditing of events may now be performed in different order. The order of these checks depends on the type of query, on whether commands were batched, and on the inclusion of procedural objects in the batch. This can affect the type of error, the error number, or the type of event audited.

## Permissions on Objects in Procedures Checked at Run-time

Users can create procedural objects (procedures, views, etc.) that name objects which they normally do not have access to, although they are not allowed to execute the procedure or use the view.

Previously, attempts to create the procedural object would fail when the create command was run. Now, the create command succeeds, but attempts to use the object fail.

For example, the following command to create a procedure succeeds, even though the user issuing the command does not have permission on mary.sales\_prospect:

```
create proc steal_leads
```

```
as
select * from mary.sales_prospects
```

However, this user cannot execute the procedure; the attempt to execute returns a permissions error.

## Stored Procedure Existence Checks in Command Batches

In pre-11.5 versions, an entire batch fails during compilation if a stored procedure named in the batch does not exist. In 11.5.x and later versions, the existence of stored procedures included in batches is not checked until execution time.

Each command in the batch succeeds or fails independently at execution time unless `begin transaction` and `commit` or `rollback transaction` commands are used, so other commands in the batch may succeed, while the stored procedure execution returns an error message.

For example, this batch attempts to execute a procedure that does not exist:

```
ANY_COMMAND_BEFORE

exec proc_that_does_not_exist

ANY_COMMAND_AFTER
```

In pre-11.5 versions, the entire batch would fail at compilation time. In 11.5.1 and later versions, the commands before and after the missing procedure execute (if they do not return errors) but `proc_that_does_not_exist` returns an error.

## Ordering of Integrity and Permissions Checks

In some cases in pre-11.5 versions, integrity checks were performed before permissions checks. Now, certain permissions checks are performed before integrity checks. This may result in cases where commands that previously failed integrity checks may instead fail permissions checks, if both these problems existed in the command. The results of auditing may differ.

## Changes in Error Numbers

Error message numbers have changed and new messages have been added. Applications that perform checks on message numbers need to also check for the new error message numbers. The permission failures described in this table include:

- Those permissions set with the grant command
- Permissions that depend on roles (such as sa\_role or SSO\_role) or on status as the DBO (database owner)

| <b>Pre-11.5.1 error number</b> | <b>Changed to:</b> | <b>Type of message</b>  |
|--------------------------------|--------------------|---|
| 174                            | 10312              | Permission failure on built-in ad hoc audit function call   |
| 229                            | 10330<br>567       | Object permission failure;<br>sp_loadbupgrade role failure<br>sp_addlanguage,<br>sp_setlangalias role failure |
| 230                            | 10332              | Column permission failure   |
| 259                            | 10321              | Ad-hoc updates to system tables not enabled   |
| 262                            | 10302, 10331       | Permission failure for setuser command  |
| 567                            | 10353<br>10314     | Permission failure for shutdown command<br>Permission failure for sp_password                                 |
| 573                            | 10317              | Password too short in sp_addlogin   |
| 574                            | 10316              | Password is the same as previous password in sp_password  |
| 916                            | 10302<br>10351     | Permission failure for dump database command<br>User is not a valid user in the database                      |
| 1915                           | 10304              | Permission failure for create index command   |
| 2557                           | 10304              | Permission failure for dbcc commands  |
| 2563, 2564, 2565, 2567         | 10302              | Permission failure for dbcc commands  |
| 2571                           | 10353              | Permission failure for dbcc commands  |
| 2710                           | 10304              | Permission failure for create, alter, truncate, update statistics, or bulk insert commands on a table         |

| <b>Pre-11.5.1 error number</b> | <b>Changed to:</b> | <b>Type of message</b>  |
|--------------------------------|--------------------|---|
| 3004                           | 10302              | Permission failure for dump database command  |
| 3110                           | 10302              | Permission failure for load database command  |
| 3704                           | 10302<br>10304     | Permission failure for drop database command<br>Permission failure for drop default command |
| 4220                           | 10302              | Permission failure for dump transaction command   |
| 4312                           | 10302              | Permission failure for load transaction command   |
| 4602                           | 10353              | Permission failure of grant create database command   |
| 4603                           | 10302              | Permission failure of grant or revoke command   |
| 4710                           | 10304              | Permission failure for truncate table command   |
| 4910                           | 10304              | Permission failure for alter table...add column command                                     |
| 4931                           | 10321              | Permission failure for alter table...add column command on a system table                   |
| 5011                           | 10302              | Permission failure for alter database command   |
| 5116                           | 10353              | Permission failure for disk mirror  |
| 6102                           | 10353              | Permission failure for kill command   |
| 7620                           | 10311              | Failure of attempt to change audit options manually.  |
| 7729                           | 11105              | Failure for set role commands; new message indicates a non-existent role                    |

| <b>Pre-11.5.1 error number</b> | <b>Changed to:</b> | <b>Type of message</b>  |
|--------------------------------|--------------------|---|
| 7730                           | 10319              | Failure for set role commands; new message indicates that the role is not granted to the user                     |
| 7731                           | 10320              | Failure for set role commands; new message indicates that the role is not granted to the user in current database |
| 7964                           | 10302              | Permission failure for dbcc commands  |
| 7983                           | 10354              | Permission failure for sp_procxmode   |
| 9503                           | 10353              | Role required message.  |
| 9902                           | 10353              | Permission failure for dbcc commands  |
| 17461                          | 567                | Permission failure for sp_remap   |

## 11.9.2 Query Processing Changes

Release 11.9.2 introduced extensive changes in query processing and optimization. Statistics became more granular and accurate. Some of these changes may result in the optimizer creating different query plans when executing old applications. These changes are discussed in the following sections:

- New and Changed Statistics
- Aggregate Optimization
- Costing Large I/O
- The optdiag Utility
- Trace Flag 302

### New and Changed Statistics

The following sections describe 11.9.2 changes to statistics, which include:

- Column Level Statistics

- Table/Index Level Statistics
- Changes to the update statistics Command

Column Level Statistics

The sysstatistics table contains statistics related to the data in a column. Column level statistics describe the distribution of values in the column (and in the case of multicolumn density values, a set of columns). These are the “distribution” statistics formerly stored in the single distribution page of each index; they include the histogram and the density values. Also stored in sysstatistics is the date and time of the last modification of the column statistics.

Column statistics are now an attribute of a column, not an index. This is a fundamental change to the statistics. Because statistics are now an attribute of a column they can be placed on any column, whether it is part of an index or not.

The number of steps (cells) to use in the histogram can be specified in the create index or update statistics commands.

Table/Index Level Statistics

The systabstats table contains statistics related to objects: the table and index level statistics. Statistics related to tables and indexes are now stored in a single table and new statistics are also available. Table and index statistics are maintained dynamically by ASE. They should not be written directly; they will be quickly overwritten. Table and index statistics cannot be written with optdiag.

The statistics stored in systabstats are dynamic. An in-memory copy is kept and modified as changes occur. This in-memory copy is then flushed to systabstats in a number of ways, such as by housekeeper, update statistics, optdiag, checkpoint, shutdown, and sp\_flushstats.

Changes to the *update statistics* Command

The update statistics command was given some new extensions and functionality. These include the ability to use update statistics to build or update statistics on an individual column, on all columns of an index or indexes or on all columns of a table. On DOL tables update statistics will use isolation level 0 (dirty reads) and avoid locking the table while running.

As in previous versions, running update statistics table\_name [index\_name] will update the statistics of the leading column (major attribute) of all indexes on the table or of the specified index. See the section Updating Statistics After Upgrade in *Chapter 6* of this manual for guidelines for running this command after upgrade.

See the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/aseperf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/aseperf/@Generic__BookView) for information on histogram cells.

## Aggregate Optimization

Aggregate optimization was introduced in ASE 11.9.x. This allows the optimizer to factor in an index on the column referenced in an aggregate function. Sometimes the optimizer can be overly optimistic in costing this index and may choose it when it is not the best choice. If you notice poor performance because of this (showplan will show the index on the aggregate column being used), you can turn off aggregate optimization with trace flag 370, which will revert the server to a pre-11.9.x handling of aggregates by the optimizer.

## Costing Large I/O

Large I/Os can now be used to access index pages along with data pages. More accurate costing of large I/Os using cluster ratios was added in 11.9.X. These new statistics values for tables and indexes are stored in systabstats.

Cluster ratios are used to determine if large I/O access will be efficient. The lower the cluster ratio, the less likely a large I/O access will be useful. Look under the heading “Derived statistics” in optdiag or trace flag 302 output (Table Information Block and the Index Selection Block) for the cluster ratios.

## The *optdiag* Utility

Release 11.9.x introduced the new command line utility optdiag. optdiag enables you to read, write and simulate statistics. Working directly with the statistics as stored in sysstatistics and systabstats is difficult because of their datatypes, optdiag makes it practical to read and modify statistics. Also, optdiag allows you to test performance based on simulated statistics.

optdiag output is essential when analyzing optimizer-related problems. It contains all the information about the data set that you need.

See the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/aseperf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/aseperf/@Generic__BookView) for information on optdiag.

## Trace Flag 302

Trace flag 302 prints the optimizer’s index costing and selection phase. It has been redesigned to provide information on the new statistics and costing in an easy-to-read format.



This trace flag is well-documented in New Functionality in 11.9.2 at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1192e/asefun92/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1192e/asefun92/@Generic__BookView) and the sections on query tuning in the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/aseperf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/aseperf/@Generic__BookView).

#### Syntax for Trace Flag 302

When using trace flag (302) to see optimizer information, use trace flag 3604 to send output to the client. (You can use trace flag 3605 to send output to the errorlog, but this is not recommended.) The syntax looks like this:

```
dbcc traceon (3604,302)
```

Trace flag 302 output has been divided into the following “blocks” of related information:

- The Table Information Block
- The Base Cost Block
- The Clause Block
- The Column Block
- The Index Selection Block
- The Best Access Block

The following is an example of trace flag 302 output:

```
*****
Beginning selection of qualifying indexes for table
'lineitem',
varno = 0, objectid 240003886.
The table (Datarows) has 600572 rows, 44308 pages,
The table's Data Page Cluster Ratio 0.999990

Table scan cost is 600572 rows, 44489 pages,
using no data prefetch (size 2K I/O),
in data cache 'default data cache' (cacheid 0) with MRU
replacement

Selecting best index for the SEARCH CLAUSE:
lineitem.l_partkey = 2000
lineitem.l_orderkey > 10000

Estimated selectivity for l_orderkey,
selectivity = 0.983249, upper limit = 1.000000.

Estimated selectivity for l_partkey,
selectivity = 0.000052, upper limit = 0.052686.
Estimating selectivity of index 'lineitem_lorder',
```

```
indid 2
scan selectivity 0.983249, filter selectivity 0.983249
590512 rows, 45991 pages, index height 2,
Data Row Cluster Ratio 1.000000,
Index Page Cluster Ratio 0.998609,
Data Page Cluster Ratio 1.000000
```

The best qualifying index is 'lineitem\_pskey' (indid 4) costing 35 pages, with an estimate of 31 rows to be returned per scan of the table, using no index prefetch (size 2K I/O) on leaf pages, in index cache 'default data cache' (cacheid 0) with LRU replacement using no data prefetch (size 2K I/O), in data cache 'default data cache' (cacheid 0) with LRU replacement

```
Search argument selectivity is 0.000051.
*****
```

See the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/aseperf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/aseperf/@Generic__BookView) for information on reading trace flag 302 output.

## If Your Version is 11.5 or 11.9.x

This section covers the following topics:

- ANSI Joins
- dsync Function
- 12.0 Query Processing Changes

### ANSI Joins

ASE now supports ANSI joins. Sybase recommends that you rewrite applications to use ANSI outer joins because they specify unambiguously whether the on or where clause contains the predicate. Transact-SQL syntax in previous releases was ambiguous in some cases.

ANSI syntax allows you to write either of these types of joins:

- Inner joins, in which the joined table includes only the rows of the inner and outer tables that meet the conditions of the on clause. The result set of a query that includes an inner join does not include any null supplied rows for the rows of the outer table that do not meet the conditions of the on clause. The syntax for an ANSI inner join is:

```
select select_list
from table1 inner join table2
on join_condition
```

For example:

```
select au_id, titles.title_id, title, price
from titleauthor inner join titles
and price > 15
```

- Outer joins, in which the joined table includes all the rows from the outer table whether or not they meet the conditions of the on clause. If a row does not meet the conditions of the on clause, values from the inner table are stored in the joined table as null values. The where clause of an ANSI outer join restricts the rows that are included in the query result. ANSI syntax also allows you to write nested outer joins. The syntax for an ANSI outer join is:

```
select select_list
from table1 {left | right} [outer] join table2
on predicate [join restriction]
```

For example:

```
select au_fname, au_lname, pub_name
from authors left join publishers
on authors.city = publishers.city
```

For more information see the Transact-SQL User's Guide at <http://manuals.sybase.com/onlinebooks/group-as/srg1100e/sqlug>.

## ***dsync* Function**

Beginning with version 12, Adaptive Server initializes UNIX file system devices using the *dsync* option to ensure recovery. This could cause a performance impact to servers using file system devices for tempdb. If you require earlier behavior or performance characteristics, change the *dsync* setting using *sp\_deviceattr*.

See the section New Unix File System Device Support in *Chapter 6* of this manual. See also the sections on disk init and sp\_deviceattr in the Reference Manual at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/refman/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/refman/@Generic__BookView) for more information.

## 12.0 Query Processing Changes

The query processing and optimizer changes introduced in ASE 12.0 are less extensive than those in 11.9.2, and should provide improved performance for most sites. These changes include the following:

- Predicate Transformation and Factoring
- 50 Tables in a Query
- Abstract Query Plans
- Increased Optimization Time
- Like Optimization Enhancements

For information on optimizer changes over recent ASE releases, see also the collection *An Introduction to Sybase Adaptive Server Enterprise's Modern Optimizer* at <http://my.sybase.com/detail?id=1013556>. This downloadable collection of optimizer-related articles published in the ISUG Technical Journal requires Acrobat Reader.

### Predicate Transformation and Factoring

Predicate Factoring provides significant performance improvement in queries with limited access paths, that is, very few possible SARGs, Joins, or or clauses that can be used to qualify rows in a table.

Additional optimization is achieved by generating new search paths based on join conditions, search clauses, and optimizable or clauses.

Predicate factoring helps query optimization by extracting optimizable clauses from predicates linked with or (difficult to optimize) and replacing them with and clauses (easier to optimize) thus providing the optimizer with more usable SARGs. The more SARGs available, the more information the optimizer has and the more likely it is to choose an efficient plan.

Full cartesian joins are avoided for some complex queries.

Example:

```
select * from lineitem, part
where ((p_partkey = l_partkey and l_quantity >= 10)
or (p_partkey = l_partkey and l_quantity <= 20) )
```

becomes

```
select * from lineitem, part
where ((p_partkey = l_partkey and l_quantity >= 10)
or (p_partkey = l_partkey and l_quantity <= 20) )
and (p_partkey = l_partkey)
and (l_quantity >= 10 or l_quantity <= 20)
```

The addition of the conjuncts adds usable SARGs for the optimizer. If new conjuncts (and clauses) added by predicate transformation and factoring are found not to be useful as part of index access strategy (that is, for filtering) they are not used.

Query semantics do not change. The result set is the same. Predicate factoring cannot be turned off. It is completely transparent to the user. It is implemented as a new compiler phase, just before the start of the optimizer.

When it happens, trace flag 302 output shows additional costing blocks and showplan shows additional “Keys are” messages.

## 50 Tables in a Query

ASE now supports up to 50 user tables and 14 worktables. This is up from 16 user tables and 12 worktables in 11.9.2. Bear in mind, however, the limits on number of referential integrity checks remain at 192 and the number of subqueries allowed remains at 16.

## Abstract Query Plans

Abstract plans were introduced in release 12.0 as a means for system administrators and performance tuners to protect the overall performance of a server from changes to query plans. Changes in query plans may be due to:

- Adaptive Server software upgrades that affect optimizer choices
- New Adaptive Server features that change query plans
- Changing tuning options such as the parallel degree, table partitioning, or indexing

Adaptive Server can capture query text and save an abstract plan for a query in a new system table called `sysqueryplans`. Using a rapid hashing method, incoming SQL queries can be compared to stored query text, and if a match is found, the saved abstract plan is used to execute the query.

For more information on abstract query plans, see the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView). For a quick overview of how to use abstract query plans to troubleshoot performance problems after upgrade, see the TechNote Using and Maintaining Abstract Query Plans at <http://my.sybase.com/detail?id=1010648>

## Increased Optimization Time

Queries with long chains of join keys may require additional time to optimize with Adaptive Server. If the time required to optimize such a query is unacceptable, consider using an abstract query plan for the query.

## Like Optimization Enhancements

ASE 12.0 changes the costing for like clauses that are not migrated into SARGs, using a technique of generating more accurate cost estimates for queries that include a leading wildcard in a like clause.

This provides better selectivity estimates, resulting in better query plans. Previously a like clause with a leading wild card was estimated to qualify all rows in the column (a selectivity of 1.0) since there was no way to search for a match. This was not the most accurate way to cost such a clause.

Example:

```
select ... from part, partsupp, lineitem
where l_partkey = p_partkey
and l_partkey = ps_partkey
and p_title = '%Topographic%'
```

The like string is compared with histogram cell boundaries.

A match occurs when the like term is found in a cell boundary by means of a pattern match search.

If no pattern match is found the selectivity is assumed to be 1/number of steps in the histogram. If the default of 20 cells is used, selectivity will be 0.05 if there is now pattern match in the histogram boundary values.

If a pattern match is found in the cell boundaries then the selectivity is estimated to be the sum of the weights of all cells with a pattern match.

In either case the resulting selectivity estimate will be more accurate than in previous versions.

This also applies to queries with like clauses of the type like `"_abc"`, or like `"[ ]abc"`.

## If Your Version is 11.5, 11.9.x, or 12.0

This section covers the following topics:

- Changes to Transact-SQL
- The `enable xact coordination` Configuration Parameter
- New Reserved Words
- Wide Columns and Data Truncation

## Changes to Transact-SQL

In general, changes to existing Transact-SQL commands add options to them and will not create a problem for applications where these changed commands appear. In some cases ASE 12.5 allows you to specify new sizes for arguments to commands. This is a result of the larger data sizes now available. If you intend to use larger data sizes, check What's New in Adaptive Server Enterprise at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic__BookView) for new size limits and other changes to Transact-SQL.

## The *enable xact coordination* Configuration Parameter

Non-DTM RPCs may fail initially on 12.0 and 12.5 when `enable xact coordination` is turned on. To prevent problems, turn this configuration parameter off with `sp_configure`.

---

**Note** This is expected to be fixed in subsequent releases of ASE 12.5.

---

## New Reserved Words

Reserved words can only be used by ASE. New ASE objects and commands result in new reserved words being added with every release. To see the new reserved words for ASE 12.0, see the Adaptive Server Reference Manual at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/aserefmn/@Generic\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/aserefmn/@Generic_BookView).

You must change all object names that contain reserved words *before you upgrade*. You must also change those names in your procedures, SQL scripts, and applications before you can run them against the upgraded server.

Use the reserved word check in `sqlupgrade`, which you can run without starting the upgrade, to check object names. Subsequent checks can be performed by running the stored procedure, `sp_checkreswords`. Names that contain reserved words can be changed with the `sp_rename` procedure, or can be enclosed in double quotes. See the discussion on reserved words in your installation guide. The following chart shows the keywords added in recent releases:



| <b>Release</b> | <b>New reserved words</b>  |
|----------------|--|
| <b>11.9.2</b>  | exp_row_size<br>reservepagegap<br>lock<br>readpast<br>reorg  |
| <b>12.0</b>    | proxy<br>proxy_table<br>key<br>jar<br>join<br>install<br>remove<br>identity_gap<br>quiesce<br>modify |
| <b>12.5</b>    | deterministic<br>func<br>function<br>inout<br>new<br>out<br>output<br>stringsize                     |

---

**Note** The Sybase procedures for detecting reserved words in object names cannot be used to find them in scripts and applications. You need to check scripts and applications separately.

---

## Wide Columns and Data Truncation

Release 12.5 introduced larger page and column sizes. This can produce unexpected results in existing applications where truncation of data used to occur. In addition, if you plan to use larger data sizes, you must take into account the ability of client software to handle the new sizes. This section gives an overview of these issues:

- Change in Truncation Behavior

- Using New Wide Data Sizes with Open Client
- Wide Columns and Optimizer Statistics
- Wide Columns with `col_length()` and `datalength()`

## Change in Truncation Behavior

Previous releases of Adaptive Server had a column-length limit of 255 bytes. The 12.5 release of Adaptive Server allows you to create columns using `char`, `varchar`, `binary`, and `varbinary` data up to 16294 bytes long, depending on the logical page size your server uses. Because of this, data that was truncated at 255 characters for earlier releases of Adaptive Server is no longer truncated. If your application depended on this truncation, the result set it receives may no longer be accurate. In the following example, `col1` and `col2` are each 200 characters long, and `col3` is 255 characters long:

```
select * from t1 where col1 + col2 = col3
```

If `col1` and `col2` each hold 200 characters, their concatenation results in a 400-character string. Previous releases of Adaptive Server truncate this to 255 characters, and the sum of `col1` and `col2` might match `col3`. However, for the 12.5 release of Adaptive Server, the product of `col1` and `col2` is 400 characters, and will never match a `col3` that is 255 characters long.

Character and binary expressions can produce a result up to 16384 bytes long. Data that is over this length is truncated.

---

**Note** See the section on New Table Limits in *Chapter 6* of this manual.

---

## Using New Wide Data Sizes with Open Client

Adaptive Server and Open Client interact depending on how you determine Open Client's capabilities. Capabilities describe features that a client/server connection supports. In particular, capabilities describe the types of requests that an application sends on a specific connection and the types of server responses that a server returns on a specific connection. Before you use Open Client with the wider limits allowed by Adaptive Server version 12.5, you must enable the Open Client capability for these limits. See your Open Client documentation for information about enabling the wide-table capabilities.

---

**Note** DB Library cannot use the wide limits for Adaptive Server version 12.5

---

To see whether you can enable wide data sizes for your client software, go to the section below that describes your system:

- Both ASE and Open Client are Version 12.5
- Open Client Is an Older Version
- ASE 12.5 Connects to a Remote Pre-12.5 Server

Both ASE and Open Client are Version 12.5

You should not have any problems if both Adaptive Server and Open Client are at version 12.5. Perform the following before you use Open Client 12.5 with Adaptive Server version 12.5:

- 1 Relink your Open Client application with the 12.5 libraries.
- 2 Establish the new version number with CS\_VERSION.
- 3 If you are using JConnect, reconnect the drivers.

Open Client Is an Older Version

If you are running a pre-version 12.5 Open Client application with Adaptive Server version 12.5, and have relinked your application with 12.5 Open Client, but have not changed the version of your libraries, the pre-version 12.5 functionality is fine, but the new limits are not enabled. Adaptive Server sends the data according to the pre-12.5 limits, and will truncate any data that is beyond these limits.

If you are running pre-version 12.5 Open Client with Adaptive Server version 12.5, but have not relinked your libraries, Adaptive Server sends data according to the pre-12.5 version, with the same restrictions outlined above.

ASE 12.5 Connects to a Remote Pre-12.5 Server

If Open Client 12.5 connects to a 12.5 Adaptive Server, which then queries a remote server which is at an earlier version, the remote server truncates any wide data before returning results. All servers and clients in a transaction must be at 12.5 in order to use the new wide limits.

## Wide Columns and Optimizer Statistics

When placing statistics on a wide column, the server uses on the first 255 bytes of data. Even if you place an index on a column created as `char(500)`, only the first 255 bytes will contribute towards the column histograms in `sysstatistics`. For `unichar` columns, the first 127 chars will be used to gather column distributions.

## **Wide Columns with `col_length()` and `datalength()`**

The commands `col_length()` and `datalength()` are “built-in” Transact-SQL functions that return database information. These functions can now return values greater than 225. If you use one of these functions in a query and assign the result to a variable, be sure to make the variable large enough to hold the value.

# Implement: Making Database Administration Changes

## In This Chapter

This chapter discusses changes to Adaptive Server system administration that can cause problems if you are not prepared for them. It does not attempt to cover new features or changes that are not likely to cause surprises. For a comprehensive listing of changes and new features, see What's New in Adaptive Server Enterprise at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic__BookView). For changes that primarily affect existing applications, see Chapter 5, “Implement: Making Required Application Changes”.

---

**Note** Changing applications and system administration is the most time-consuming part of migration preparation and this guide does not tell you how to make these changes. You must choose a method to do this that suits your needs and resources. For instance, you may wish to develop your own plans and scripts, or you may prefer to contact Sybase Consulting for help. Sybase Technical Support offers a Web site, the ASE Migration Resources Web page at <http://sybase.com/support/techdocs/migration>, for migration information and can help with technical problems such as bugs and error conditions.

---

Begin reading at the earliest section that includes your current ASE version:

- If Your Version is 11.5
- If Your Version is 11.5 or 11.9.x
- If Your Version is 11.5, 11.9.x, or 12.0

See also the white paper Technical Issues in ASE 11.9.x and 12.0 Upgrade at <http://sybase.com/detail/1,3693,1011767,00.html> for some platform-specific issues affecting ASE 12.x.

## If Your Version is 11.5

This section covers the following topics:

- New Rollback Records
- Optimizer Changes
- Trace Flag 326

### New Rollback Records

New rollback records added to the transaction log in 11.9.2 may require you to add more space to the syslogs system table. Rollback records are logged whenever a transaction is rolled back. For every update record that is rolled back, a rollback record is logged.

To make sure that there is always space for rollback records in the event of a rollback, ASE saves enough space to log a rollback record for every update belonging to an open transaction. If a transaction completes successfully, no rollback records are logged and the space reserved for them is released.

To calculate the increased amount of space that needs to be added to the transaction log to accommodate rollback records, you need to estimate:

- The number of update records in the transaction log that are likely to belong to already rolled-back transactions.
- The maximum number of update records in the transaction log that are likely to belong to open transactions at any one time.

Each rollback record requires approximately 60 bytes of space, or 3 one hundredths of a 2K page. Thus, the calculation for including rollback records (RRs) in the transaction log is:

$$(\text{logged RRs} + \# \text{ open updates}) * 3/100$$

In addition to increasing the size of the transaction log to handle rollback records, it may also be useful to add log space to compensate for the effects of rollback records on the last-chance threshold (LCT) and on user-defined thresholds.

## Optimizer Changes

ASE 11.9.2 introduced major changes in the way the server keeps statistics and optimizes queries. These changes were aimed at increasing granularity in the statistics kept by the server, improving query plans, and producing better performance. Changes introduced in 11.9.2 included:

- Two new system tables, `systabstats`, which maintains information on objects such as tables and rows, and `sysstatistics` which stores data distribution statistics.
- New optimizer statistics and added functionality to the `update statistics` command.
- The `optdiag` utility, which simplifies the task of looking at optimizer statistics. This utility also allows you to model query performance by running queries against a test database with simulated table, index, and column statistics.

These and other query processing changes for 11.9.2 are reviewed in the chapter Chapter 5, “Implement: Making Required Application Changes” and in the section *Summary of Query Processing Changes in 11.9.2 and 12.0 that May Affect Performance* in the white paper Technical Issues in ASE 11.9.x and 12.0 Upgrade at <http://sybase.com/detail/1,3693,1011767,00.html>.

For more detailed information on the `systabstats` and `sysstatistics` tables, see the chapter on “Using Statistics to Improve Performance” in the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView). For information on the `optdiag` utility, see the Utility Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/util/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/util/@Generic__BookView)

## Updating Statistics After Upgrade

As part of your upgrade procedure, you need to run:

```
update statistics table_name
```

on all tables as soon as possible following upgrade. Statistics that are copied from the old distribution page by the upgrade are not as accurate as those created by `update statistics`, because the weights are average, not actual. Running this command updates default statistics on the leading columns of your indexes.

However, be cautious about adding new statistics to the production environment until you have determined your needs and tested any changes. Not all new statistics are needed or helpful in every environment.

---

**Warning!** Making unneeded changes to statistics may adversely affect performance. In addition, running the `update statistics` command with incorrect parameters may change the number of steps in a histogram and impact performance. Apply changes first in a test environment. See the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView) for information.

---

## Trace Flag 326

trace 326 no longer exists in ASE 11.9.2 or above. In pre-11.9.2 trace 326 used arithmetic averaging to establish the total density value rather than the default geometric averaging. If you used trace 326 in earlier versions of ASE (perhaps in the RUNSERVER file with the -T flag) you now have to manually change the total density value via `optdiag`.

## If Your Version is 11.5 or 11.9.x

This section covers the following topics:

- Documentation Changes
- Backup Server
- Suspending Updates to Databases for Copying
- New Unix File System Device Support
- Changes to dbids
- CIS Proxy Databases
- ASE Plug-in for Sybase Central
- The sybsyntax Utility



## Documentation Changes

ASE 12.0 includes several new manuals that describe separately-licensed Adaptive Server features such as Java, High Availability, and Distributed Transaction Management.

### Information Moved to Different Manual

Information from the following manuals:

- *Security Features User's Guide*
- *Security Administration Guide*

has been moved to the following manuals:

- *System Administration Guide*
- *Performance and Tuning Guide*

Information from *Configuring Adaptive Server for UNIX* has been moved into the installation guide.

### Discontinued Documents

The ASE documentation set no longer includes the following documents:

- *Adaptive Server Enterprise Getting Started*
- *Introducing Adaptive Server Enterprise*
- *Master Index to Adaptive Server Documentation*
- *Navigating Adaptive Server Enterprise Documentation*
- *Online sybsyntax database and sp\_syntax procedure*
- *Quick Reference Card*

## Backup Server

This section is to alert you to changes in Backup Server that you may find advantageous. You are not required to make any changes in your current maintenance routines if you do not wish to. To find out more about Backup Server, see the System Administration Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/asesag/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/asesag/@Generic__BookView).

## Stripes

In earlier versions, Backup Server could dump to a maximum of 32 stripes, either disk or tape. In 12.0, more stripes are available. The maximum number of stripes it can use is equal to the maximum number of Open Server threads it can create.

---

**Note** The documentation for ASE 12.0 still refers to 32 stripes, but this is an error.

---

## Network Connections

The maximum number of network connections a local Backup Server can originate is limited by Open Server to 9118. Therefore, the maximum number of remote stripes that Backup Server can use in a single dump or load operation is 9118.

## File Descriptors

Backup Server uses two file descriptors for each stripe apart from the file descriptors associated with the error log file, interfaces file and other system files. However, there is a per-thread limitation imposed by the operating system on the number of file descriptors. Open Server has a limitation of 1280 on the number of file descriptors that an application can keep track of.

## Non-Rewinding Tapes

In ASE 12, the tape head is positioned at the end of valid dump data, which saves time when you want to perform multiple dump operations. In previous versions of ASE, Backup Server rewound the tape before each dump operation.

## Starting Backup Server

When you start Backup Server, use the `-P` flag to configure the Open Server for the maximum number of threads it will create. The maximum number of threads equals the maximum number of stripes available. If you have started Backup Server without setting a high enough `-P` value and you attempt to dump or load a database to a number of stripes that exceeds the number of threads, the dump or load operation will fail.

You must configure the local and remote Backup Servers at startup by providing the appropriate values to the command line options. A remote dump to greater than 25 stripes with the local and remote Backup Servers started with default configuration will fail because the maximum number of network connections that Backup Server can originate (specified by the -N option) is by default 25. The maximum number of server connections into the remote Backup Server (specified by the -C option) is 30.

## Dump Format

In Backup Server 12.0 the stripe number must be stored in integer format. Earlier versions of Backup Server stored the stripe number in the HDR1 label in ASCII format in 4 bytes. Thus, earlier Backup Server versions will not be able to load a dump file that uses the version 12.0 dump format. However, Backup Server 12.0 can read and write earlier versions of the dump format. Dump or load operations involving one or more remote servers will abort with an error message in the following instance: if one or more remote servers uses an earlier Backup Server version and the number of stripes to which the database is dumped (or from which it is loaded) is greater than 32.

## Suspending Updates to Databases for Copying

Adaptive Server version 12 provides the new quiesce database command to temporarily suspend updates to databases. You may want to suspend database updates to use an external database backup utility or to unmirror a disk. While the database is in the quiescent state, read transactions from the database are allowed. You can suspend and restart transactions with the quiesce database command issued from different user sessions.

If you have multiple databases on a single device, and you want to copy the entire device, you need to quiesce all the databases on the device before you perform the copy. If you have multiple devices for a single database, be sure to copy all the devices.

For more information on the quiesce database command, see the System Administration Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView) for ASE 12.5.

## New Unix File System Device Support

In release 12.0, Sybase introduced the `dsync` option of the `disk init` command. This option guarantees that writes to database devices placed on Unix file systems are performed by the operating system as soon as the server flushes the data. In the past, the server might log a write, but the operating system buffered the data. If the system went down, the server had no way to recover the buffered data that was lost.

The `dsync` option could cause a performance problem in some cases. If you place `tempdb` on a file system, you should consider disabling `dsync` on that device. For more information about using the `dsync` option to protect data on Unix file systems, see the System Administration Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView).

## Changes to *dbids*

In ASE 12, the `dbids` of the following databases start at 31513:

- `dbccalt`
- `dbccdb`
- `sybsecurity`
- `sybssystemdb`
- `sybssystemprocs`

If you drop and recreate these databases after an upgrade, ASE applies these new `dbids`. The `dbids` for all other databases are determined in the same way as in earlier versions.

---

**Note** If you perform an upgrade in place, that is, you upgrade without rebuilding your system, you should not see any change in `dbid` between versions. This is an issue only if your maintenance scripts use these IDs and you build a new system from scratch, which will then use the new ID conventions.

---

## CIS Proxy Databases

In 12.0, HA creates databases with names like “\_\_pxy”. As a result, you cannot drop proxy databases with similar names ending in “pxy” from Adaptive Servers configured with Sybase Failover in a high availability system. If you attempt to drop these proxy databases, Adaptive Server issues an error message. This is expected behavior.

If you need to drop a database with this type of name, change configuration so your system is not configured for high availability.

## ASE Plug-in for Sybase Central

The Adaptive Server plug-in to Sybase Central no longer displays all servers listed in the *sql.ini* file. Instead, Sybase Central lists only those servers that you connected to earlier, or those servers that are started as Windows NT services. To access a new server for the first time, use the Connect menu option to select a server listed in the *sql.ini* file.

---

**Note** You need to have Java installed to run the Sybase Central.

---

## The *sybsyntax* Utility

The *sybsyntax* utility, which provides online help with Transact-SQL syntax, was removed from ASE 12.0. However, *sybsyntax* was part of ASE 12.5. Unless you are upgrading to ASE 12.0, you do not need to concern yourself with this change.

## If Your Version is 11.5, 11.9.x, or 12.0

This section covers the following topics:

- Changes to Documentation
- System Database Sizes
- Changes to System Tables
- New Database and Table Limits

- New Reserved Words
- Configuration Parameters
- Increased Memory
- New Storage for text and image Data

## Changes to Documentation

The following changes were made to the Adaptive Server documentation set:

- *Managing and Monitoring Adaptive Server Enterprise* has been discontinued.
- The *Utility Guide* is now a generic book that includes all utilities, regardless of platform.

## System Database Sizes

When you upgrade a server, the following system databases require more space than may be required by a new installation. Be sure that you have around 10% free space for upgrade.

The installation guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic\\_\\_CollectionView;pt=asp1250e](http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e) gives the minimum space requirements for system databases in systems using 2k or larger page size.

## The *sybsystemprocs* Database

You may need to increase the size of the *sybsystemprocs* database and possibly the device it is on. See your installation guide for a discussion of enlarging *sybsystemprocs*.

If *sybsystemprocs* is placed on multiple device fragments, all fragments need to have both data and log segments. This means that if the *sysusages* table has multiple rows for *sybsystemprocs*, the *segmap* column value must be “7” for every row.

## **sybssystemdb**

The sybssystemdb database is used for the spt\_values table and two-phase commits. Before ASE 12.5, this database was not required. The database is automatically created on the master device by the configuration utility.

Because this database may have a lot of activity and use a lot of log space, we strongly recommend that you create sybssystemdb on another device prior to upgrade or move it to another device after upgrade. This protects the master device in case the sybssystemdb log fills up.

If you will not use the features that require this database, you can keep it at the minimum size.

## **Changes to System Tables**

There are no new system tables in ASE 12.5. However, two types of changes could cause trouble if you do not adjust your scripts to account for them:

- Additional columns in some system tables
- Change to larger datatypes in some columns

If you have scripts that refer to system tables, make appropriate changes.

## **BCP with syslogins**

If you try to bcp data into the syslogins system table from an earlier version of the server, bcp will fail. This is because in 12.0 and 12.5, two additional columns were added to syslogins. If you need to bcp into syslogins, you can work around the problem by creating a dummy table to hold the data, altering the table to add two columns, and then copying the data into the 12.5 syslogins.

---

**Note** This issue could occur with other system tables that have added columns. You might encounter the problem first with syslogins if you need to copy this data as part of the upgrade.

---

See What's New in Adaptive Server Enterprise at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic__BookView) for changes to system tables in ASE 12.0 and 12.5.

## New Database and Table Limits

This section covers changes in table and database limits:

- Maximum Number of Users and Logins
- New Table Limits

### Maximum Number of Users and Logins

The maximum number of logins to the server and users to the database has been increased. ASE 12.5 now supports:

- 2 billion plus 32K logins to the dataserver
- 2 billion less 1032193 users to a database
- 1032193 groups to a database

There is some new behavior due to this change:

- Negative values may be used for user IDs (uid)
- The server user ID (suid) associated with a group or a role in sysusers is no longer equal to the negation of their uid. In ASE 12, every suid associated with a group or a role in sysusers is set to -2(INVALID\_SUID).

See What's New in Adaptive Server Enterprise at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic__BookView) for ID ranges.

### New Table Limits

The following chart shows new table object limits for the different page sizes, with some differences according to locking scheme. This is for your information only. For more information on creating tables, see the Transact-SQL User's Guide at <http://manuals.sybase.com/onlinebooks/group-as/srg1100e/sqlug>

| User-visible object (sizes in bytes) | 2K logical page    | 4K logical page    | 8K logical page    | 16K logical page    |
|--------------------------------------|--------------------|--------------------|--------------------|---------------------|
| Number of rows per page              | APL=256<br>DOL=166 | APL=256<br>DOL=337 | APL=256<br>DOL=678 | APL=256<br>DOL=1361 |



| User-visible object (sizes in bytes) | 2K logical page      | 4K logical page      | 8K logical page      | 16K logical page       |
|--------------------------------------|----------------------|----------------------|----------------------|------------------------|
| Row length                           | APL=1960<br>DOL=1958 | APL=4008<br>DOL=4006 | APL=8104<br>DOL=8102 | APL=16296<br>DOL=16294 |
| Fixed-length column                  | APL=1960<br>DOL=1958 | APL=4008<br>DOL=4006 | APL=8104<br>DOL=8102 | APL=16296<br>DOL=16294 |
| Variable-length column               | APL=1947<br>DOL=1952 | APL=3987<br>DOL=4000 | APL=8067<br>DOL=8096 | APL=16227<br>DOL=16288 |
| Size of index key                    | APL=600<br>DOL=600   | APL=1250<br>DOL=1250 | APL=2600<br>DOL=2600 | APL=5300<br>DOL=5300   |

## New Reserved Words

Reserved words, or ASE keywords, can only be used by ASE and cannot be used to name user-created objects. New ASE objects and commands result in new reserved words being added with every release. The following chart shows the keywords added in ASE 11.9.2, 12.0, and 12.5:

| Release | New reserved words   |
|---------|--|
| 11.9.2  | exp_row_size<br>reservepagegap<br>lock<br>readpast<br>reorg  |
| 12.0    | proxy<br>proxy_table<br>key<br>jar<br>join<br>install<br>remove<br>identity_gap<br>quiesce<br>modify |

| Release | New reserved words  |
|---------|---|
| 12.5    | deterministic<br>func<br>function<br>inout<br>new<br>out<br>output<br>stringize |

You must change all object names that contain reserved words *before you upgrade*. You must also change those names in your procedures, SQL scripts, and applications before you can run them against the upgraded server.

Use the reserved word check in `sqlupgrade`, which you can run without starting the upgrade, to check object names. Subsequent checks can be performed by running the stored procedure, `sp_checkreswords`. Names that contain reserved words can be changed with the `sp_rename` procedure, or can be enclosed in double quotes. See the discussion on reserved words in your installation guide.

---

**Note** The Sybase procedures for detecting reserved words in object names cannot be used to find them in scripts and applications. You need to check scripts and applications separately.

---

## Configuration Parameters

Configuration parameter values that you have set are preserved by the upgrade. Only those parameters that are set at “DEFAULT” in the server configuration file are changed to the new version’s default value. Because some parameter default values are now higher, values that you set in the older version may not be high enough for the new version. The following parameters can cause problems during or after upgrade if they are set too low:

- stack size
- cpu grace time
- enable housekeeper GC

The parameter `enable xact coordination` can cause a problem for RPCs when set at default. These issues are explained the next sections.

## stack size

The default value of the stack size parameter has increased over the last several releases. While stack size varies from platform to platform, the changes on Solaris (shown in the table below) illustrate the increase in stack size requirements with successive versions:

| Server Version | Minimum Stack Size                           |
|----------------|--|
| 11.0.x         | 24576 (24K)                                  |
| 11.5           | 34816 (32k)                                  |
| 11.9.2         | 34816 (32k)                                  |
| 12.0           | 46090 (45K) - 32-bit<br>86016 (84K) - 64-bit |
| 12.5           | 46090 (45K) - 32-bit<br>86016 (84K) - 64-bit |

If the stack size is set too low, the server could experience stack overflow, either during upgrade or later during processing.

Check your stack size parameter and reset it to default if necessary. You may do this by modifying the server configuration file, replacing the existing value with the word “DEFAULT”.

**Note** ASE on Solaris may require a stack size two to four times the default in environments with heavy I/O using file system devices.

For more information on configuration parameters, see the System Administration Guide at <http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/asesag/>. For a comprehensive list of new and changed configuration parameters, see What's New in Adaptive Server Enterprise at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/whatsnew/@Generic__BookView).

## RPC Problem with “enable xact coordination”

The configuration parameter `enable xact coordination` defaults to “1”. With this parameter turned on, you cannot execute an RPC from within a transaction. Set this parameter to “0” if you are not using DTM.

## Increased Memory

ASE 12.5 uses more memory than previous releases. ASE now allocates more memory for the dataserver, procedure cache, and named caches. If you do not increase physical and/or shared memory, the additional needed memory will be taken from the default data cache, possibly affecting performance. Check the installation guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic\\_\\_CollectionView;pt=asp1250e](http://manuals.sybase.com:80/onlinebooks/group-as/asp1250e/@Generic__CollectionView;pt=asp1250e) for minimum memory requirements. Use `sp_helpcache` to check cache sizes before and after upgrade. Adjust memory if necessary to restore the original value of the default data cache after upgrade.

See the System Administration Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/sag/@Generic__BookView) for information on configuring memory and data caches; see the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView) for more information on memory issues.

## New Storage for *text* and *image* Data

As part of 12.0 performance enhancements, Sybase replaced the old scan method used to access text and image data with an access method based on a new st-node structure. which wholly or partially (depending on the size of the text data) resides on the first text page. The st-node contains pointers to each set of contiguous text pages. For a description of how st-nodes work, see the TechNote New Text and Image Data Storage at <http://my.sybase.com/detail?id=1014269>.

Because of problems arising in early versions of ASE 12.0 and 12.5, it is recommended that you apply some bug fixes before accessing your text and image pages. See the following TechNotes for a discussion of text/image issues.

- Loss of text/image Index Information in Upgraded Databases Could Lead to Access Errors in ASE 12.0 and 12.5 at <http://my.sybase.com/detail?id=1012553>
- Problems Affecting text data Following Upgrade to Adaptive Server Enterprise 12.0 or 12.5 at <http://my.sybase.com/detail?id=1012153>
- Avoiding text and image Problems With ASE 12.0 and 12.5 at <http://my.sybase.com/detail?id=1013522>

## ***buildmaster* Command Obsolete**

ASE 12.5 no longer uses the *buildmaster* command. Instead, the *dataserver* command or *sqlsrvr.exe* for NT (the ASE product binary) has a new flag, *-b*, to run in build mode. Use the *dataserver -b* command where you used *buildmaster* in the past.

Because of changes in the master device and database, as well as the elimination of the *buildmaster* command, recovery is more complex if you lose your master device or database following an upgrade. Before upgrading, make both a physical dump and a *bcp* copy of the master database.

See the Utility Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/util/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/util/@Generic__BookView) for more information on the *dataserver* command. See the Troubleshooting and Error Messages Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/svrtsg/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/svrtsg/@Generic__BookView) for information about disaster recovery.



# Test: Ensuring Stability and Performance

## In This Chapter

This chapter will help you evaluate testing methods and develop a testing plan. It contains these sections:

- The Goal of Testing
- Setting Up the Test Environment
- Prioritizing Applications to be Tested
- Establishing Performance Criteria
- Developing Fallback Procedures
- Summary of Testing Techniques
- Writing Performance Scripts
- Test Cycle: Summary of Tests
- Test Cycle: Testing for Performance

---

**Note** This chapter is intended as a guide to various kinds of tests. You should use those suggestions that fit your business needs and resources. Full-scale testing as described here is highly recommended but not required for migration success.

---

## The Goal of Testing

The primary goal of testing is to ensure that after migration:

- Application behavior is predictable.
- Application and operational service levels are preserved or exceeded.

- The test and production systems are stable and the data is safe.
- The upgrade is successful and does not adversely impact the production system.

## **Setting Up the Test Environment**

Ideally, you should set up a dedicated hardware configuration (including subnets) and Adaptive Server exactly like your production system. Creating an identical system lets you make valid comparison, perform real tuning as part of migration effort, and if you wish to do so, switch the test system to production later on. Create the test system as described in the sections below:

- Make Backups
- Use Scripts to Create the Test System
- Create Your Databases by Loading Backups
- If the Test Environment Is Not an Exact Duplicate

### **Make Backups**

Make backups of the production system. You can use these to populate your test system and to restore it when necessary.

### **Use Scripts to Create the Test System**

Using the object creation scripts you gathered, wrote or reverse engineered in Chapter 3, “Analyze: Documenting Your Environment”, build a test environment matching your production system.



Use backups or the bcp scripts to populate your test databases.

---

**Note** When you create a new database and then load bcp files, you reduce the fragmentation that may be present in the production system and change the performance characteristics of the database. If this is not desirable because you do not intend to rebuild your production environment, you may prefer to create your databases as described in the next section, Create Your Databases by Loading Backups.

---

Run dbcc's in the test databases immediately to be sure that there are no problems at this stage.

## Create Your Databases by Loading Backups

If you do not intend to rebuild your production environment, you may prefer to create your test databases with the for load option of the create database command. This makes your test databases more representative of the current production environment in terms of fragmentation and density.

Follow these steps:

- 1 Create the database with the for load option of the create database command.
- 2 Load the backups you made of the production database.
- 3 Issue the online database command. This command automatically upgrades the database if it is not at the 12.5 level.

For command syntax, see the Reference Manual at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/refman/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/refman/@Generic__BookView).

## If the Test Environment Is Not an Exact Duplicate

If you have to use a smaller system for testing, try to have components identical to the production system, such as operating system level, drivers, and disk types.

You need to make adjustments when you have less disk space or memory in the test system. Scale down databases proportionally, while retaining the same data distributions so optimizer decisions remain constant. Scale down memory to ensure consistent I/O rates.

Reproduce data layout across available devices as closely as possible.

If you use fewer CPUs, adjust transaction arrival rates (that is, load) and concurrent users proportionally.

## Prioritizing Applications to be Tested

Since it may be difficult to test all application functions, gather user input to identify the most critical transactions. Write tests for those functions using the techniques described in the following sections. Functions not tested at this time can be validated during user acceptance tests.

## Establishing Performance Criteria

Depending on the migration plan you developed, you may want to meet or exceed the performance you get with your current system. For example, you might decide to apply guidelines like these:

- For *parallel with replication*:
  - Be sure to measure the overhead of the replication mechanism. For example, if the replication costs 10%, do not begin parallel operations until release 12.5 is tuned to outperform by 10%.
  - For an around-the-clock operation, you may decide that an initial goal of breaking even is reasonable.
- For *cutover without replication*, gear the migration for equivalent performance between the old and new systems. A goal of breaking even the first week is reasonable.

- A *phased cutover* is subject to the highest performance expectations. Some performance tuning of the production workload after cutover of the production server may be best. You can time the production server cutover to occur as soon as performance gains are acceptable and testing is successful.

See Chapter 4, “Plan: Writing a Plan and Getting Ready to Migrate” for information on migration methods.

## Developing Fallback Procedures

Before you begin testing, be sure that you know how you will return the test system to a known state when you have a problem. You will use the same fallback plan when you migrate the production system.

During simple testing cycles, it’s sometimes faster to back out unwanted changes. However, this is not recommended for *timed runs* during benchmarking. You have to restore from backup after each timed run to return the system to a known state.

Back up all databases before and after the test system upgrade, as you would for a “real” upgrade. The backups preserve the layout of data on disk and help you avoid confusion due to fragmentation and page splits.

To ensure source code control, use scripts for all changes to objects. This makes it much easier to recreate your environment if necessary.

## Summary of Testing Techniques

You can use a variety of testing techniques in your test plan. The following table summarizes the advantages and disadvantages of various testing techniques and tools:

| Technique                            | Description   | Advantages   | Disadvantages  |
|--------------------------------------|---|--|--|
| Ad hoc testing                       | Manually walk through important application processes, screens, and reports | <ul style="list-style-type: none"> <li>• Easy to implement</li> <li>• Tests front-end applications and back-end servers</li> </ul>   | <ul style="list-style-type: none"> <li>• For complex applications, code coverage is too small</li> <li>• Difficult to distinguish front-end and back-end bottlenecks if response time is a determining factor</li> <li>• Impossible to obtain production multiuser load, which misses concurrency and capacity issues altogether</li> </ul>  |
| Manual performance scripts and cases | Specify input and compare with known outputs                                | <ul style="list-style-type: none"> <li>• Easy to implement</li> <li>• Common basis for regression test suites</li> <li>• Back-end focus may help locate the cause of a problem</li> </ul>  | <ul style="list-style-type: none"> <li>• Only tests back- end server</li> <li>• Impossible to obtain production multiuser load, which misses concurrency and capacity issues altogether</li> <li>• No ad hoc query testing</li> <li>• Depends on strong analysis of process or transaction profiles</li> </ul>   |
| Keystroke capture                    | Record and replay keystrokes and mouse clicks into an application           | <ul style="list-style-type: none"> <li>• Tests front-end applications and back-end servers</li> <li>• Tool may include powerful language and looping capabilities to manipulate inputs for multiuser concurrency and capacity testing</li> </ul> | <ul style="list-style-type: none"> <li>• Heavy processing requirements - may require additional hardware</li> <li>• May increase development time for creating multiuser test simulations, and add time for debugging test harness</li> <li>• Depends on strong analysis of process or transaction profiles</li> </ul>   |
| Concurrency and capacity testing     | Use third-party load-testing tools  | <ul style="list-style-type: none"> <li>• Tests both front and back ends</li> <li>• Language and looping constructs make these incredibly powerful tools</li> </ul>   | <ul style="list-style-type: none"> <li>• Heavy processing requirements for test tool—may even require additional hardware or else results could be skewed</li> <li>• Learning and development curve to write multi-user test simulations</li> <li>• Risk of bugs in test harness could skew results</li> <li>• Dependent on strong analysis of process/transaction profiles</li> </ul> |

| Technique               | Description   | Advantages   | Disadvantages   |
|-------------------------|---|--|---|
| Transaction generation  | Thin client to simulate user execution of transactions  | <ul style="list-style-type: none"> <li>• Strong multiuser load testing</li> <li>• Focus on back-end server issues</li> </ul>   | <ul style="list-style-type: none"> <li>• May increase development time for creating multiuser test simulations, though learning and development curve generally less than keystroke capture tools</li> <li>• Adds time for debugging test harness to prevent skewed results</li> <li>• Depends on strong analysis of process or transaction profiles</li> </ul> |
| Production load capture | Using tools to capture real transactions in a production environment, including performance and semantic characteristics, and resubmit in a test environment for analysis | <ul style="list-style-type: none"> <li>• Tests real production loads, including ad hoc queries</li> <li>• Especially useful when little or no analysis of transaction profiles is available</li> </ul> | <ul style="list-style-type: none"> <li>• Introduces new software into a production environment</li> <li>• Production and test system configurations must be identical for valid performance analysis</li> </ul>   |

## Writing Performance Scripts

This section discusses the basics of writing performance scripts:

- Write Benchmark Scripts
- Drivers

### Write Benchmark Scripts

In general, you have to write special benchmark scripts rather than rewriting applications as benchmarks.

To write the benchmark script:

- Add a function to `funcs.c` for each transaction
- Generate any run-time data required (such as primary key to select or data to insert)
- Write code to submit SQL or stored procedure(s) to Adaptive Server (for example, using `dbsqlexec()` calls)

- Write code to process result set(s)
- Name each transaction explicitly (for example, “begin tran cust\_update”) to make it easier to identify in system procedures and tables
- For a stored procedure-based system, verify the parameters that make the stored procedures work. If the parameters need to vary for a meaningful test, add the necessary logic.

Volume is critical in performance simulation. A script roughly equivalent to an application, running at the normal production volume for that application, is usually better than a script functionally matched to an application, but running at only half the volume.

If workload is based on client PCs issuing Transact-SQL, you can use performance monitoring tools, available in the market, to capture data streams.

## Drivers

In addition, you need drivers for these tasks:

- General Error Handling
- Deadlock Handling
- Result Handling
- Time Measurement
- Runtime Data Generation

### General Error Handling

In the event of an error, you can either throw away the transaction and not count it or, depending on the requirements of the test, restart the transaction and count the entire response time.

### Deadlock Handling

You have to resubmit deadlocked transactions to get a realistic time measurement. You can count average response time with and without deadlocks.

## Result Handling

You can capture query results by fetching the entire result set back to the client and binding it to local variables. Outputting results to a file may increase your time by requiring additional file I/O and operating system buffering.

## Time Measurement

Time measurement is especially important in multi-tiered applications where a bottleneck could occur at any level. Generally, you should time database transactions, not business functions. Start at transaction “send” and stop at the last row processed in the result set. Logical business operations can be aggregated later.

Granularity is important for problem identification and resolution. Important measurements to take include:

- Throughput (number of transactions per second/minute/hour)
- Average and maximum response time by transaction
- A histogram of response time ranges by transaction. For example:
  - Less than 1 second
  - Between 1 and 2 seconds
  - Between 2 and 3 seconds
  - Over 3 seconds

## Runtime Data Generation

When using runtime-generated data, skewed keys may be a problem. Selects may show an unrealistically high cache hit rate, and inserts, updates, and deletes may appear to have concurrency problems. To avoid these problems, use the entire range of key values in your tests, unless you are trying to recreate a particular business situation.

Using a separate file for each transaction can prevent skewing; however, it requires additional code to synchronize users. Using a separate file for each user is more work to create but also prevents skewing.

Using a memory generator works well to prevent skewing and makes the test easy to administer. However, the benchmark is then not 100% repeatable.

## Test Cycle: Summary of Tests

This section summarizes a complete test cycle, with tests that target specific issues, including old and new functionality, performance under multi-user loads, integration, and user acceptance.

| Stage                         | Purpose   | Best Technique   |
|-------------------------------|---|--|
| Functional testing            | <p>For each application or process, addresses the following questions:</p> <ul style="list-style-type: none"> <li>• Are there any obvious bugs?</li> <li>• Do transactions return the same results?</li> <li>• Will the application break anywhere?</li> </ul> <p>If you decide to use new functionality:</p> <ul style="list-style-type: none"> <li>• Does the release provide the functionality it markets?</li> <li>• What are the limitations of the new features of this release?</li> </ul> | <p>Single-user:</p> <ul style="list-style-type: none"> <li>• Ad hoc test</li> <li>• Manual test scripts and cases</li> <li>• Existing application test suites</li> </ul> |
| Stress testing (benchmarking) | <p>Using very heavy loads, addresses the following questions:</p> <ul style="list-style-type: none"> <li>• Are there any bugs related to multi-user loads?</li> <li>• Is the performance of critical transactions as good or better?</li> <li>• Is the new release stable under load?</li> </ul>  | <p>Multiuser:</p> <ul style="list-style-type: none"> <li>• Keystroke capture</li> <li>• Transaction generator</li> <li>• Production load capture</li> </ul>              |
| Integration testing           | <p>Ensure that all system components work well together, such as:</p> <ul style="list-style-type: none"> <li>• Batch processing</li> <li>• Online transaction processing (OLTP)</li> <li>• Decision Support Systems (DSS) and ad hoc queries</li> <li>• Operations, including backup, recovery, and dbcc commands</li> <li>• Sybase products other than ASE</li> <li>• Third party products</li> </ul>  | <p>Test suite models all system components</p>   |
| End-user acceptance testing   | <p>Execute acceptance tests specific to the environment. Also cover functions not prioritized into earlier stages.</p> <p><b>Note</b> The other stages, done well, should have caught most of the problems.</p>   | <p>Standard acceptance tests</p>   |



| Stage                        | Purpose  | Best Technique   |
|------------------------------|--|--|
| Final migration plan testing | <ul style="list-style-type: none"> <li>• Ensure that you are fully prepared by walking through the upgrade/migration plan.</li> <li>• Verify that fallback procedures work.</li> <li>• Identify and test the contingencies.</li> </ul> | Walk through every upgrade step, including fallback strategies |

## Test Cycle: Testing for Performance

This section expands on performance benchmarking before and after upgrade. For a summary of the entire test cycle, see, “Test Cycle: Summary of Tests” on page 10.

This section covers the following tasks:

- Pre-Upgrade Single-User Tests
- Pre-Upgrade Multi-User Tests
- Test System Upgrade
- Post-Upgrade Single-User Tests
- Post-Upgrade Multi-User Tests

---

**Note** Benchmarks test server processing speed. Therefore, run your benchmarks from a light-weight client so that all processing happens at the back-end.

---

### Pre-Upgrade Single-User Tests

After you create the test system at the same ASE level as the production system and before running benchmarks, you need to sync the test system with the production system so that comparisons to the new server are valid.

### Optimizer

Compare showplan and output from dbcc 310 and 302 which you gathered in Chapter 2, “Analyze: Documenting Business Requirements”, to that of the test system to verify that the optimizer is making the same decisions. Some possible causes of optimizer differences include:

- Data and data distribution problems. Check data layout and be sure that these match the production system.
- Index definitions. Verify that your reindexing scripts worked. Using `sp_helpindex`, compare the index definitions of the production and test systems.

Resolve any problems before going on.

## I/O

Verify that the test server is performing the same amount of I/O. Physical and logical I/O should be in the same proportion as in the production system. Use `statistics io` output to compare.

## Pre-Upgrade Multi-User Tests

When the optimizer is behaving the same as it did in the production system, you are ready to run the following benchmarks.

- Untimed Benchmarks
- Timed Benchmarks

### Untimed Benchmarks

In multi-user mode, run untimed benchmarks as often as necessary, capturing response time and throughput metrics. Resolve:

- Differences, such as cache hit rate, to the production system
- Bottlenecks caused by saturated devices
- Other problems or mistakes

Between runs, restore your databases from backup or perform a quick reset.

### Timed Benchmarks

Always restore your databases from backup so they are in a known state when running timed benchmarks. Gather response time and throughput metrics. Run the tests as often as necessary to fix problems and be able to reproduce the final results.

## Test System Upgrade

Perform the upgrade on the test system following the instructions in the installation guide. Be sure to perform all steps so that the test upgrade is a walkthrough of the real upgrade. Make the memory and disk space changes you determined that you needed in Chapter 6, “Implement: Making Database Administration Changes”. However, do not change your Adaptive Server configuration to use new features at this time.

**Note** Your first objective after upgrade is to test Adaptive Server’s out-of-the-box performance. Wait until later to try new performance tuning features.

Resolve all problems and note them for reference. The installation guide gives detailed troubleshooting information.

After the upgrade:

- Perform dbccs on all databases.
- Back up all databases, including master and sybsystemprocs. You will need backups to restore the test system to a known state between timed benchmarks.

## Post-Upgrade Single-User Tests

In single-user tests, verify that:

- The optimizer is creating the same or better query plans. Differences in query plan may be due to optimizer changes, depending on the ASE version you were running. Determine whether you need to change application code to compensate. See Chapter 7, “Test: Ensuring Stability and Performance”, for possible optimizer issues.
- I/O counts, physical and logical, are the same as they were before the upgrade.

## Post-Upgrade Multi-User Tests

After reloading backups to return the test system to a known state, run the following benchmarks.

- Untimed Benchmarks
- Timed Benchmarks

## Untimed Benchmarks

In multi-user mode, run untimed benchmarks as often as necessary, capturing response time and throughput metrics. Resolve:

- Differences, such as cache hit rate, to the pre-upgrade state
- Bottlenecks
- Other problems or mistakes

Be sure that you have enough default data cache after the upgrade. Tune as needed to fix problems. Between runs, restore your databases from backup or perform a quick reset.

Refer to the Performance and Tuning Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic\\_\\_BookView](http://manuals.sybase.com:80/onlinebooks/group-as/asg1250e/perf/@Generic__BookView) for more information. See also Technical Information Library for the latest TechNotes and White papers on issues that may affect performance.

## Timed Benchmarks

Always restore your databases from backup so they are in a known state when running timed benchmarks. Gather response time and throughput metrics.

Tune as necessary to achieve your performance criteria.

Run tests as often as necessary to fix problems and be able to reproduce the final results.

# Worksheets for Your Current Environment

This appendix provides guidelines and sample worksheets for gathering information that will help you in planning for migration. The topics are:

- Adaptive Server Operational Worksheets
- Data Architecture Worksheet
- Adaptive Server Infrastructure Worksheets

As part of migration planning, see the ASE Migration Resources Web page at <http://sybase.com/support/techdocs/migration> for current documents.

## Adaptive Server Operational Worksheets

The following worksheets are for gathering business requirements. For more information, see Chapter 2, “Analyze: Documenting Business Requirements”.

This worksheet includes the following sections:

- Operational Business Requirements
- Backup and Restore Procedures
- Database Dump Details
- Maintenance Procedure Details

## Operational Business Requirements

The following worksheet is for recording operational business requirements:



| Task  | Yes/No | Comments |
|---|--------|----------|
| Are backup and recovery procedures documented?          |        |          |
| Are automated dump procedures in place?                 |        |          |
| What is the number of generations of dumps?             |        |          |
| Are dumps kept offsite?                                 |        |          |
| Are maintenance activities documented?                  |        |          |
| What is the frequency of ASE error log scanning?        |        |          |
| Is database space utilization monitored?                |        |          |
| Has database capacity planning been performed recently? |        |          |

## Database Dump Details

Detailed backup and restore information is helpful for defining success criteria:

| Database Name | Frequency of Database Dumps | Dump Device Used | Frequency of Transaction Log Dumps | Dump Device Used | Comments |
|---------------|-----------------------------|------------------|------------------------------------|------------------|----------|
|               |                             |                  |                                    |                  |          |

| Database Name | Frequency of Database Dumps | Dump Device Used | Frequency of Transaction Log Dumps | Dump Device Used | Comments |
|---------------|-----------------------------|------------------|------------------------------------|------------------|----------|
|               |                             |                  |                                    |                  |          |
|               |                             |                  |                                    |                  |          |
|               |                             |                  |                                    |                  |          |
|               |                             |                  |                                    |                  |          |
|               |                             |                  |                                    |                  |          |
|               |                             |                  |                                    |                  |          |
|               |                             |                  |                                    |                  |          |
|               |                             |                  |                                    |                  |          |
|               |                             |                  |                                    |                  |          |
|               |                             |                  |                                    |                  |          |

### Maintenance Procedure Details

Use the following worksheet to record detailed maintenance information:

| Database/Object Name | Frequency of <i>dbcc checkdb/checktable</i> | Frequency of <i>dbcc checkalloc/tablealloc</i> | Frequency of <i>update statistics</i> | Frequency of Monitoring Space Utilization |
|----------------------|---|--|---------------------------------------|---|
|                      |   |  |                                       |   |
|                      |   |  |                                       |   |
|                      |   |  |                                       |   |
|                      |   |  |                                       |   |
|                      |   |  |                                       |   |







- Measure I/Os per second per disk and controller, and I/O queue lengths per “time window” per server.
- Also measure total I/Os, reads, and writes per second per Sybase device per “time window” per server.
- Concurrency
 

Determine the average lock contention. You can use `sp_who` to determine the processes currently running, and `sp_lock` to find out which current processes hold locks.
- Network I/O
  - Measure the packets per second per NIC per “time window” per server.
  - Also measure the TDS packets (“sent from” and “received by”) per “time window” per server.
- Memory
  - Determine the paging/swapping rates per second per “time window” per server.
  - Also determine the data and procedure cache hit rates per “time window” per server.

## Transaction Profile

Record transaction profile information for each database on the server:

---

| Process<br>(Xact)<br>Name | Process<br>Type<br>(OLTP,<br>batch) | Xact<br>Priority | Frequ<br>ncy per<br>User<br>(per<br>hour) | Source<br>Code Type<br>(stored<br>procedure,<br>ESQL, etc.) | Average<br>Respon<br>se Time<br>Required | Maximu<br>m<br>Respon<br>se Time<br>Required | Current<br>Average<br>and<br>Maximu<br>m<br>Respon<br>se Time |
|---------------------------|-------------------------------------|------------------|---|---|--|--|---|
|---------------------------|-------------------------------------|------------------|---|---|--|--|---|

---



Record this information for both the *production* and *development* environments.

## Host Configuration

This section provides worksheets for host configuration.

### Hardware

Record technical support information about your hardware manufacturer:

---

**ASE Machine**

---

Make:

---

Model:

---

Customer ID with hardware vendor:

---

Technical support phone number:

---

Technical support hours:

---

Technical account manager:

- Name
  - Phone number (s)
  - Pager number
- 

Record CPU resources:

---

**CPU**

---

Number of physical processors:

---

Chip speed:

---

Number of processors available to ASE:

---

---

**CPU**

---

Other CPU-intensive processes/threads:

---

Processes/threads bound to specific CPUs:

---

Processes/threads run at high priority:

---

**Physical Memory Usage**

List all the major processes and memory requirements running on each server.

---

| <b>Application</b> | <b>Runtime Memory Usage</b> |
|--------------------|-----------------------------|
| Operating system   |                             |

---

ASE total memory + additional netmem + extent i/o buffers

| <b>Application</b>   | <b>Runtime Memory Usage</b> |
|--|-----------------------------|
| Open Servers   |                             |
| Include:   |                             |
| <ul style="list-style-type: none"> <li>• Backup Server</li> <li>• sybmultbuf</li> <br/> <li>• Replication Server</li> <li>• Monitor Server</li> <br/> <li>• Gateways (list)</li> </ul> |                             |
| Other applications   |                             |
| <i>Total memory required</i>   |                             |
| <i>Total memory installed</i>  |                             |

### Disk I/O Configuration

General disk information can help with firmware incompatibilities and capacity planning.

| <b>Controller Number</b> | <b>Make and Model</b> | <b>Firmware Revision</b> | <b>Months in Service</b> | <b>Transfer Rate (KB/sec)</b> |
|--------------------------|-----------------------|--------------------------|--------------------------|-------------------------------|
|                          |                       |                          |                          |                               |
|                          |                       |                          |                          |                               |
|                          |                       |                          |                          |                               |
|                          |                       |                          |                          |                               |

The following disk layout information can help with firmware incompatibilities, nearing Mean-time Between Failures (MTBF), load balancing, and capacity planning.

| Physical Device Name | Make and Model | Firmware Revision | Months in Service | Controller Number | Capacity (MB) | Throughput (I/Os per second) | Transfer Rate (KB per second) |
|----------------------|----------------|-------------------|-------------------|-------------------|---------------|------------------------------|-------------------------------|
|                      |                |                   |                   |                   |               |                              |                               |
|                      |                |                   |                   |                   |               |                              |                               |
|                      |                |                   |                   |                   |               |                              |                               |
|                      |                |                   |                   |                   |               |                              |                               |
|                      |                |                   |                   |                   |               |                              |                               |
|                      |                |                   |                   |                   |               |                              |                               |
|                      |                |                   |                   |                   |               |                              |                               |
|                      |                |                   |                   |                   |               |                              |                               |

The following disk layout information can help in case redistribution is required, for load balancing, and for capacity planning.

| Physical Device Name | Partition Number | Used by (Sybase, UNIX File System, etc.) | Device Name | OS Mirrored Device Name | Capacity (MB) | Cylinder Range |
|----------------------|------------------|--|-------------|-------------------------|---------------|----------------|
|                      |                  |  |             |                         |               |                |
|                      |                  |  |             |                         |               |                |
|                      |                  |  |             |                         |               |                |
|                      |                  |  |             |                         |               |                |
|                      |                  |  |             |                         |               |                |
|                      |                  |  |             |                         |               |                |
|                      |                  |  |             |                         |               |                |
|                      |                  |  |             |                         |               |                |

The following logical volume information can help in case redistribution is required, for load balancing, and for capacity planning.



| Logical Volume Device | Member Disk Partitions | Used by (Sybase, UNIX File System, etc.) | Sybase Device | Mirror Logical Device | Capacity (MB) | Stripe Width (MB) |
|-----------------------|------------------------|--|---------------|-----------------------|---------------|-------------------|
|                       |                        |  |               |                       |               |                   |
|                       |                        |  |               |                       |               |                   |
|                       |                        |  |               |                       |               |                   |
|                       |                        |  |               |                       |               |                   |
|                       |                        |  |               |                       |               |                   |
|                       |                        |  |               |                       |               |                   |
|                       |                        |  |               |                       |               |                   |
|                       |                        |  |               |                       |               |                   |

### Network Configuration

Network layout information can help with firmware incompatibilities, MTBF, and capacity planning.

| Physical Device Name | Make and Model | Firmware Revision | Months in Service | Supported Protocols | Network Address | Transfer Rate (KB/second) |
|----------------------|----------------|-------------------|-------------------|---------------------|-----------------|---------------------------|
|                      |                |                   |                   |                     |                 |                           |
|                      |                |                   |                   |                     |                 |                           |
|                      |                |                   |                   |                     |                 |                           |
|                      |                |                   |                   |                     |                 |                           |
|                      |                |                   |                   |                     |                 |                           |
|                      |                |                   |                   |                     |                 |                           |
|                      |                |                   |                   |                     |                 |                           |
|                      |                |                   |                   |                     |                 |                           |

## Tape Configuration

Tape layout information can help with firmware incompatibilities, MTBF, and capacity planning.

---

| Physical Device Name | Make and Model | Firmware Revision | Months in Service | Capacity (MB) | Controller Number | Transfer Rate (KB/second) |
|----------------------|----------------|-------------------|-------------------|---------------|-------------------|---------------------------|
|                      |                |                   |                   |               |                   |                           |
|                      |                |                   |                   |               |                   |                           |
|                      |                |                   |                   |               |                   |                           |
|                      |                |                   |                   |               |                   |                           |
|                      |                |                   |                   |               |                   |                           |

---

## Operating System Configuration

Detail the operating system.

---

### Server Hardware

Name:

---

Release level:

---

Patch history:

---

Kernel configuration parameters:

---

Swap space size:

---

Technical support phone number:

---

Technical support hours:

---

---

**Server Hardware**

---

Technical account manager:

- Name
  - Phone number(s)
  - Pager number
- 

## Adaptive Server Configuration

Document general information about the ASE configuration.

---

**General Configuration**

---

Home directory:

---

Components, release and fix levels:

---

---

Locations/names of scripts to rebuild database environment:

---

---

sp\_configure configuration values:

---

---

buildmaster configuration values:

---

Run dbcc memusage on ASE during an off-peak time or in single-user mode.

| Usage                  | MB | 2K Pages | Bytes |
|------------------------|----|----------|-------|
| Configured memory      |    |          |       |
| Code size              |    |          |       |
| Kernel structures      |    |          |       |
| Server structures      |    |          |       |
| Page cache             |    |          |       |
| Proc buffers           |    |          |       |
| Proc headers           |    |          |       |
| Number of page buffers |    |          |       |
| Number of proc buffers |    |          |       |

## Database Devices

Database device information can help in case redistribution is required, load balancing, and capacity planning.

| Database Device Name | Physical Device Name | Virtual Device Number | Capacity (2K Pages) | Mirrored Device Name |
|----------------------|----------------------|-----------------------|---------------------|----------------------|
|                      |                      |                       |                     |                      |
|                      |                      |                       |                     |                      |

| Database Device Name | Physical Device Name | Virtual Device Number | Capacity (2K Pages) | Mirrored Device Name |
|----------------------|----------------------|-----------------------|---------------------|----------------------|
|                      |                      |                       |                     |                      |
|                      |                      |                       |                     |                      |
|                      |                      |                       |                     |                      |

## Databases and Segments

Database and segment information can help with load balancing and capacity planning.

| Database Name | Database Options Set | Fragment (page range) | Size (in MB) | Segment Names | Device Name |
|---------------|----------------------|-----------------------|--------------|---------------|-------------|
|               |                      |                       |              |               |             |
|               |                      |                       |              |               |             |
|               |                      |                       |              |               |             |
|               |                      |                       |              |               |             |
|               |                      |                       |              |               |             |
|               |                      |                       |              |               |             |
|               |                      |                       |              |               |             |
|               |                      |                       |              |               |             |

## Dump Devices

Dump device information can help with load balancing and capacity planning.

| Database Device Name | Physical Device Name | Media Type | Capacity (MB) |
|----------------------|----------------------|------------|---------------|
|                      |                      |            |               |
|                      |                      |            |               |
|                      |                      |            |               |
|                      |                      |            |               |
|                      |                      |            |               |
|                      |                      |            |               |
|                      |                      |            |               |

# Sample Migration Task Lists

This appendix provides the following samples:

- Sample Task List Template
- General Migration Task List Example
- Parallel Migration Task List Example
- Cutover Migration Task List Example
- Staged Cutover Task Overview

---

**Note** These examples are illustrations. They are not step-by-step procedures that must be followed by everyone. Every migration is unique and you need to write a plan appropriate for your site.

---

## Sample Task List Template

The following table is a suggested format for detailing the tasks you need to perform as part of your migration. Modify it as necessary:

| Task | Date | Begin Time | End Time | Duration | Owner | Status |
|------|------|------------|----------|----------|-------|--------|
|      |      |            |          |          |       |        |
|      |      |            |          |          |       |        |
|      |      |            |          |          |       |        |

## General Migration Task List Example

The following general task list shows typical migration tasks, with differing details for the install/load approach (install an new server, then load compatible dumps from earlier server) and one using the upgrade approach (run the upgrade program on the current databases). Your own migration task lists may differ in detail and order. The following general areas are covered:

- Migration Analysis
- Migration Preparation
- Implement Migration (Using Install/Load Technique)
- Implement Migration (Using Upgrade Technique)
- Migration Quality Assurance

See Chapter 4, “Plan: Writing a Plan and Getting Ready to Migrate” for more information on migration planning.

## Migration Analysis

### Document Current Configuration

- 1 Establish cutoff point for environment.
- 2 Document current server installation information.
- 3 Document current server configuration values.
- 4 Document hardware configuration.
- 5 Document applications per server.
- 6 Document application server requirements.
- 7 Document application client requirements.
- 8 Document related software and middleware.
- 9 Retrieve source database creation scripts.
- 10 Retrieve source object creation scripts.
- 11 Get counts of all server and database objects.
- 12 Review configuration document.



- 13 Update configuration document.

### **Gather Business Requirements**

- 1 Define business requirements.
- 2 Define constraints.
- 3 Define application dependencies.
- 4 Define dataserwer dependencies.
- 5 Prioritize applications.
- 6 Identify vendor issues.
- 7 Review requirements document.
- 8 Update requirements document.

### **Conduct Compatibility Analysis**

- 1 Analyze hardware compatibility.
- 2 Analyze operating system compatibility.
- 3 Analyze other Sybase software compatibility.
- 4 Analyze non-Sybase software compatibility.
- 5 Analyze middleware/API compatibility.
- 6 Analyze communications compatibility.
- 7 Analyze client platforms compatibility.
- 8 Document analysis results.
- 9 Review compatibility analysis.
- 10 Update compatibility analysis.

### **Develop Migration Strategy**

- 1 Draft migration strategy.
- 2 Review strategy (team).
- 3 Update strategy (team).
- 4 Review strategy (user/sponsor).

- 5 Define migration downtime impact.
- 6 Notify affected departments.
- 7 Revise implementation plan.
- 8 Obtain user signoff for migration and implementation.

## Migration Preparation

### Write Test Plans and Test Scripts

- 1 Write system functional tests.
- 2 Write integration tests.
- 3 Write stress tests.
- 4 Write user acceptance tests.
- 5 Review test plans.
- 6 Update test plans.
- 7 Create test scripts for each phase of testing.
- 8 Execute scripts on source system to establish baseline.
- 9 Get user signoff on test plans and baseline results.

### Prepare Applications For Migration

- 1 Search for new reserved words.
- 2 Search for subqueries.
- 3 Verify new database conversion and computations.
- 4 Search for queries containing changed SQL syntax.
- 5 Design code changes.
- 6 Get user signoff on application changes.

### Design and Develop Server Migration Scripts

- 1 Design and develop server configuration scripts.

- 2 Design and develop file system configuration for devices.
- 3 Create database device scripts.
- 4 Prepare security, login, and password corrections.
- 5 Create security scripts.

### **Design and Develop Database Migration Scripts**

- 1 Create database scripts.
- 2 Create database object creation scripts.
- 3 Create database security scripts.
- 4 Modify/create system administration scripts.

### **Design and Develop Data Migration Scripts**

- 1 Design and develop data extract scripts (such as bcp).
- 2 Determine optimal bulk copy options.
- 3 Design and develop data load scripts.

### **Perform Other Pre-migration Tasks**

- 1 Design fallback tasks.
- 2 Obtain user signoff on fallback tasks.
- 3 Perform backups.
- 4 Set up source code control environment.
- 5 Set up new user environment.
- 6 Develop other migration aids.

## **Implement Migration (Using Install/Load Technique)**

### **Create Target Environment**

- 1 Verify target system readiness.
- 2 Move migration scripts to target system.

- 3 Configure file system.
- 4 Complete installation/upgrade preparation from installation guide.
- 5 Install Adaptive Server, Backup Server, and Open Client.

### **Perform Server Migration**

- 1 Restrict access to systems and servers.
- 2 Create database devices.
- 3 Execute server security and other scripts.

### **Perform Database Migration**

- 1 Execute database creation scripts
- 2 Create database partitions and segments.
- 3 Create database objects (defaults, constraints, rules, views, etc.).
- 4 Create database tables and stored procedures.
- 5 Execute database security script.

### **Perform Data Migration**

- 1 Execute source data extraction scripts.
- 2 Move data from source to target platform.
- 3 Execute data load scripts.

### **Complete Server and Data Migration**

- 1 Create indexes and triggers.
- 2 Run dbcc commands.
- 3 Dump databases.

### **Perform Application Migration**

- 1 Develop application code changes.
- 2 Unit test application changes.

- 3 Analyze and correct application changes.
- 4 Configure applications to access new server.
- 5 Perform preliminary tuning.
- 6 Allow user access to system and server.
- 7 Insure no access to obsolete files.
- 8 Notify users of availability.

## **Implement Migration (Using Upgrade Technique)**

### **Upgrade Adaptive Server**

- 1 Verify target system readiness.
- 2 Configure file system.
- 3 Install software.
- 4 Use sqlupgrade to perform upgrade.

### **Complete Migration**

- 1 Run dbcc commands.
- 2 Dump databases.

### **Perform Application Migration**

- 1 Develop application code changes.
- 2 Unit test application changes.
- 3 Analyze and correct application changes.
- 4 Configure applications to access new server.
- 5 Perform preliminary tuning.
- 6 Allow user access to system and server.
- 7 Insure no access to obsolete files.
- 8 Notify users of availability.

## Migration Quality Assurance

### Perform System Tests

- 1 Perform functional tests.
- 2 Compare functional test results to baseline.
- 3 Analyze functional test results.
- 4 Perform corrective actions.
- 5 Retest as necessary.
- 6 Document functional test results.
- 7 Obtain user signoff on functional test results.

### Perform Integration Tests

- 1 Perform integration tests.
- 2 Compare integration test results to baseline.
- 3 Analyze integration test results.
- 4 Perform corrective actions.
- 5 Retest as necessary.
- 6 Document integration test results.
- 7 Obtain user signoff on integration test results.

### Perform Stress Tests

- 1 Perform performance/capacity tests.
- 2 Compare capacity test results to baseline.
- 3 Analyze capacity test results.
- 4 Perform corrective actions.
- 5 Retest as necessary.
- 6 Document capacity test results.
- 7 Obtain user signoff on capacity test results.

## Perform User Acceptance Tests

- 1 Perform acceptance tests.
- 2 Compare acceptance test results to baseline.
- 3 Analyze acceptance test results.
- 4 Perform corrective actions.
- 5 Retest as necessary.
- 6 Document acceptance test results.
- 7 Obtain user signoff on acceptance test results.

## Perform Production Data Refresh

- 1 Schedule production data cutover.
- 2 Extract production data.
- 3 Move data to target.
- 4 Load data into Adaptive Server production environment.
- 5 Backup new environment.
- 6 Notify users of data refresh.

## Parallel Migration Task List Example

This sample task list covers a parallel migration with replication approach to upgrade an older system to Adaptive Server Enterprise 12.5.

The scenario for this task list is:

- A major telecommunications company cannot afford system downtime for the upgrade.
- In the event of failure, the system cannot be down more than 15 minutes and the company cannot afford more than 1 hour of data loss. The migration must support these contingency requirements if a production back-out is required.

This example does not repeat the preparation steps to given in the previous example. The task lists for this example include:

- Define Test/Acceptance Criteria—Regression Test Suites
- Set Up Target Production Environment
- Set Up Replication Server
- Run Regression Test Suites
- Upgrade Server B (Shadow)
- Run Post-upgrade Regression Test Suites on ASE 12.5 (Server B)
- Run User Acceptance Tests on ASE 12.5 (Server B)
- Shift Production Users to Adaptive Server 12.5 (Server B)
- Perform Final Steps

## **Define Test/Acceptance Criteria—Regression Test Suites**

### **Back-end regression test suite—production loads**

- 1 Identify back-end queries.
- 2 Encapsulate back-end queries.
- 3 Create back-end test suite (showplan, stat io, and stat time wrappers).

### **Front-end simulation regression test suite**

- 1 Identify target user functions.
- 2 Capture and map SQL for target user functions.
- 3 Encapsulate SQL for user functions.
- 4 Create front-end simulated test suite (showplan, stat io, and stat time wrappers).

### **Front-end Phase 1 and 2 regression test suite**

- 1 Identify front-end test scenarios.
- 2 Review front-end application and acceptance/test procedures.
- 3 Document functional test approach.
- 4 Compose front-end test mix matrix.



## Set Up Target Production Environment

- 1 Identify physical drive configuration for Server A (production) and Server B (shadow).
- 2 Configure physical drives.
- 3 Perform a dump of the production environment.
- 4 Install older system on Server B.
- 5 Configure older system on Server B, to duplicate Server A.
- 6 Update the interfaces.
- 7 Create the databases.
- 8 Load the Server A dump on Server B.
- 9 Run dbcc commands (checktable, checkalloc, checkcatalog).
- 10 Synchronize user IDs.
- 11 Run checkpoint.

## Set Up Replication Server

- 1 Install Replication Server on Server B.
- 2 Configure Replication Server on Server B.
- 3 Install Replication Server on Server A.
- 4 Configure Replication Server on Server A (secondary).
- 5 Verify replication functionality between the two servers.
- 6 Test replication on target objects.
- 7 Verify and checkpoint replication-ready environment.
- 8 Recycle server.

## Run Regression Test Suites

### Back-end regression test suite—production loads

- 1 Update back-end scripts.

- 2 Iteratively perform back-end regression testing.
- 3 Monitor and capture system dynamics (sp\_who, sp\_lock, statistics io...).
- 4 Verify and document older system regression test.

### **Front-end simulation regression test suite**

- 1 Iteratively do front-end simulation regression tests.
- 2 Monitor and capture system dynamics (sp\_who, sp\_lock...).
- 3 Verify and document older system regression test.

### **Front-end Phase 1 and 2 regression test suite**

- 1 Perform Phase 1 and 2 regression testing for local team.
- 2 Perform Phase 1 and 2 regression testing for users.
- 3 Monitor and capture dynamics (sp\_who, sp\_lock...).
- 4 Make “go/no go” call for the shadow server.
- 5 Verify and document older system regression test.
- 6 Verify performance and functions on Server B.

### **Upgrade Server B (Shadow)**

- 1 Alter sybsystemprocs.
- 2 Perform pre-upgrade verification.
- 3 Upgrade Server B to ASE 12.5.
- 4 Load ASE 12.5 environment from media dumps.
- 5 Run dbcc commands (checktable, checkalloc, checkcatalog).
- 6 Set baseline 12.5 configuration parameter values.
- 7 Perform dump of older databases from Server A.
- 8 Load dump into 12.5 system.
- 9 Run dbcc commands (checktable, checkalloc, checkcatalog) on ASE 12.5 databases. Be sure to verify the dbcc log and error log.

10 Recycle ASE.

## **Run Post-upgrade Regression Test Suites on ASE 12.5 (Server B)**

### **Back-end regression test suite—production loads**

- 1 Perform back-end regression testing.
- 2 Monitor and capture system dynamics (sp\_who, sp\_lock, statistics io...).
- 3 Verify and document regression test.

### **Front-end simulation regression test suite**

- 1 Perform front-end simulation regression testing.
- 2 Monitor and capture system dynamics (sp\_who, sp\_lock...).
- 3 Verify and document regression test.

### **Front-End Phase 1 and 2 Regression Test Suite**

- 1 Perform Phase 1 and 2 regression testing.
- 2 Monitor and capture system dynamics (sp\_who, sp\_lock...).
- 3 Verify and document regression test.

### **Other Testing**

Verify 12.5 performance and functions.

## **Run User Acceptance Tests on ASE 12.5 (Server B)**

- 1 Production testers perform user regression testing of Release 12.5.
- 2 Monitor and capture system dynamics (sp\_who, sp\_lock...).
- 3 Verify and document regression test.
- 4 Determine if Release 12.5 is “go/no go”.

## Shift Production Users to Adaptive Server 12.5 (Server B)

- 1 Ensure that there is no production activity.
- 2 Perform a dump from Server A (production).
- 3 Load the dump onto Server B (shadow).
- 4 Run dbcc commands (checktable, checkalloc, checkcatalog) on the older release databases just loaded to Server B. Be sure to verify the dbcc log and error log
- 5 Switch the IP address, and rename the machines and servers.
- 6 Run user testing and verification.

## Perform Final Steps

- 1 Enable replication (Server B to Server A).
- 2 Start production users on Release 12.5 (Server B).

## Cutover Migration Task List Example

This sample task list covers a cutover migration approach. The scenario for this task list is:

- The mid-sized company requires a simple, somewhat fault-tolerant upgrade.
- In the event of failure, the company depends on nightly backups. The system cannot be down more than one hour and cannot afford more than eight hours of data loss.
- The environment includes development, test, and production systems.

This example does not repeat the preparation steps given in previous examples. This section includes the following task lists:

- Set Up ASE 12.5 Environment on Development System
- Define Test/Acceptance Criteria—Regression Test Suites
- Define Fallback Procedures on Test System
- “Baseline” Older Environment on Test System

- Run Regression Test Suites on Older Release Test System
- Upgrade Test System to Release 12.5
- Run Regression Test Suites on Release 12.5 Test System
- Run User/Acceptance Tests on Release 12.5 Test System
- Execute Fallback Procedures on Test System
- Upgrade Production Server to ASE 12.5
- Perform Final Steps

## **Set Up ASE 12.5 Environment on Development System**

- 1 Configure the physical drives and local array, duplicate the environment of the current development system.
- 2 Perform a dump of the current development system.
- 3 Install ASE 12.5.
- 4 Configure ASE 12.5 on the development system, duplicating the current development environment.
- 5 Update the interfaces file.
- 6 Create the databases on ASE 12.5.
- 7 Load the dump of the current development system.
- 8 Run dbcc commands (checktable, checkalloc, checkcatalog) on ASE 12.5 databases. Be sure to check the dbcc log and error log.
- 9 Synchronize user IDs between the old and new development systems.
- 10 Run checkpoint on the 12.5 environment.
- 11 Shift development to ASE 12.5. Developers begin verification and new feature development.

## **Define Test/Acceptance Criteria—Regression Test Suites**

### **Front-end simulation regression test suite**

- 1 Identify target user functions.

- 2 Capture and map SQL for target user functions.
- 3 Encapsulate SQL for user functions.
- 4 Create front-end simulated test suite (showplan, stat io, and stat time wrappers).

### **Front-end regression test suite**

- 1 Identify front-end test scenarios.
- 2 Understand front-end application and acceptance/test procedures.
- 3 Document functional test approach.
- 4 Compose front-end test mix matrix.

### **Define Fallback Procedures on Test System**

- 1 Create the fallback scripts to re-create the older environment.
- 2 Document fallback procedures.

### **“Baseline” Older Environment on Test System**

Recycle ASE.

### **Run Regression Test Suites on Older Release Test System**

#### **Front-end simulation regression test suite**

- 1 Iteratively perform front-end simulation regression testing.
- 2 Monitor and capture system dynamics (sp\_who, sp\_lock...).
- 3 Verify and document the regression test.

#### **Front-end Regression Test Suite**

- 1 Perform regression testing for local team.
- 2 Monitor and capture system dynamics (sp\_who, sp\_lock...).

- 3 Make the “go/no go” call for the test system.
- 4 Verify and document the regression test.
- 5 Verify performance and functions on the older release test server.

## **Upgrade Test System to Release 12.5**

- 1 Perform a dump of the current system’s databases.
- 2 Before the upgrade, run dbcc commands (checktable, checkalloc, checkcatalog).
- 3 Alter the sybssystemprocs database.
- 4 Perform a pre-upgrade verification.
- 5 Upgrade the test system to Release 12.5.
- 6 Run dbcc commands (checktable, checkalloc, checkcatalog) on the Release 12.5 databases. Be sure to verify the dbcc log and error log.
- 7 “Baseline” the Release 12.5 configuration.

## **Run Regression Test Suites on Release 12.5 Test System**

Recycle server before starting tests.

### **Back-End Regression Test Suite—Production Loads**

- 1 Iteratively perform the back-end regression testing.
- 2 Monitor and capture system dynamics (sp\_who, sp\_lock, statistics io...).
- 3 Verify and document regression test.

### **Front-End Simulation Regression Test Suite**

- 1 Perform front-end simulation regression testing.
- 2 Monitor and capture system dynamics (sp\_who, sp\_lock...).
- 3 Make the “go/no go” call for the 12.5 system.
- 4 Verify and document this regression test.

## Front-End Regression Test Suite

- 1 Perform user regression testing.
- 2 Monitor and capture system dynamics (sp\_who, sp\_lock...).
- 3 Verify and document this regression test.

## Other Testing

Verify 12.5 performance and functions on the test system.

## Run User/Acceptance Tests on Release 12.5 Test System

- 1 Testers perform user regression testing of Release 12.5.
- 2 Monitor and capture system dynamics (sp\_who, sp\_lock...).
- 3 Verify and document the regression test.
- 4 Make the Release 12.5 “go/no go” decision for the production system.

## Execute Fallback Procedures on Test System

- 1 Shut down the Release 12.5 test system.
- 2 Recreate the older (current production) release test system.
- 3 Recreate the older environment, using your recreation scripts and procedures.
- 4 Load the older production system dumps.
- 5 Run dbcc commands (checktable, checkalloc, checkcatalog) on the older databases. Be sure to verify the dbcc log and error log.
- 6 Run checkpoint on the older release test system.

## Upgrade Production Server to ASE 12.5

- 1 Dump the databases on the current production system.
- 2 Before the upgrade, run dbcc commands (checktable, checkalloc, checkcatalog).



- 3 Alter the sybssystemprocs database.
- 4 Perform a pre-upgrade verification.
- 5 Upgrade the production system to Release 12.5.
- 6 Run dbcc commands (checktable, checkalloc, checkcatalog) on the Release 12.5 databases. Be sure to verify the dbcc log and error log.
- 7 “Baseline” the Release 12.5 configuration on the production system.
- 8 Perform user testing and verification.

## Perform Final Steps

- 1 Start production users on the Release 12.5 production system.
- 2 Upgrade the test system to Release 12.5.

## Staged Cutover Task Overview

This sample task overview covers a staged cutover migration approach to rebuild Adaptive Server. The scenario for this task list is:

- The large company has a mature client/server environment with multiple applications residing on a single server.
- In the event of failure, the company depends on nightly backups and transactions dumps. The system cannot be down more than one hour and cannot afford more than two hours of data loss.
- The environment includes development, test, and production platforms. It requires enough space for duplicate servers on each platform.
- The object-level rebuild strategy necessitates moving the data, which requires system downtime for the target application.

## Tasks

The high-level tasks in this scenario are:

- 1 Configure ASE 12.5 on the development system, duplicating the earlier release development configuration.
- 2 Migrate the development objects to the 12.5 development system.
- 3 Construct regression test suites.
- 4 Construct bcp scripts to move object deltas.
- 5 “Baseline” the older test environment.
- 6 Configure a duplicate 12.5 server on the test platform.
- 7 Migrate the test/acceptance objects to the 12.5 test system.
- 8 Conduct regression test suites on the test system.
- 9 Verify synchronization of objects between the old and new release test systems.
- 10 Conduct user acceptance testing on the 12.5 test system.
- 11 Configure a duplicate 12.5 server on the production platform.
- 12 Migrate the production objects to the 12.5 production system.
- 13 Move the production users to the 12.5 production system.
- 14 Use your bcp scripts to resynchronize objects between the old and new release production systems.

## Migration Issue Checklists

The following checklists were produced by Sybase Consulting to use in writing migration plans for customers. They list issues that you may wish to address in your migration planning. Not all issues shown here will apply to your site, nor are these lists exhaustive. For more information on writing a migration plan, see Chapter 4, “Plan: Writing a Plan and Getting Ready to Migrate”.

### Logical Data Architecture

Does the logical data architecture include:

- A graphical representation of the data model from the logical architecture
- A usable map (like a CRUD matrix) of how data is used by the organization
- An approach for how data will be used by the organization
- An approach for how data will be distributed across platforms, locations
- An approach for how replicated data will be maintained
- An approach for how legacy data will be synchronized
- An approach for how data stores will be accessed, updated and purged

## Logical Application Architecture

Does the logical application architecture include:

- A list and brief description of the required RPCs and stored procedures needed to support the new IT architecture
- List of any potentially sharable functions
- List of shared services such as initialization, termination, global edits, error handling, logging, and monitoring
- An approach and a graphical representation of how application functionality will be split between servers, clients and middleware
- List of functional controls that need to be in place in the new architecture, for example audit procedures
- A description of how the new IT architecture applications will integrate with legacy applications
- Any graphical user interface standards that the new IT architecture applications must comply with
- A list of services that the new IT architecture applications must interface with, such as image, voice-mail, email, word processing, printing, fax or file transfer mechanisms
- An approach for how the performance expectations for the new IT architecture applications will be met
- An approach for how the availability requirements for the new IT architecture applications will be met

## Logical Technology Architecture

Does the logical technology architecture include:

- Vendor-independent description (functions and features) and graphical representation of hardware and software components in the new systems infrastructure
- Information about expected network loads at normal and peak processing periods
- Information about any upgrades necessary to the existing network infrastructure, including information on specific carrier types
- Information on all connected nodes (workstations, database servers, gateways, etc.), including:
  - Projected quantities
  - Role in the new IT architecture (client, service provider etc.)
  - Platform characteristics for protocol handling, storage capacities, performance, fault tolerance and security
- Information on and graphical representation of the geographical location of the system's infrastructure components
- An approach for how the availability requirements will be met in terms of hardware/software capabilities, like fault tolerance, hot standbys, etc.

## **Logical Support Architecture**

Does the logical support architecture include:

- Information on systems management procedures to be upgraded or put in place for:
  - Software distribution
  - Performance and fault monitoring
  - Fault management

- Disaster recovery
- Production approval and access control
- New support organization, including roles and responsibilities
- Staffing and training plans
- Strategy for meeting needs for support coverage (location and number of shifts)
- Strategy for meeting needs for required response time for problem resolution

## Migration Strategy Design

Does the migration strategy include:

- An implementation sequence or a transition plan for candidate applications, that shows the evolution of the new IT architecture
- Information on key migration constraints, such as, data conversions and critical interfaces to legacy systems
- Information on bridging routines to keep legacy systems synchronized while the new IT architecture is evolving
- Information on methodologies, techniques and tools to be used in application development environments
- Initial strategy for putting new IT architecture applications into production
- Staffing, skill and training requirements to construct and test new IT architecture applications
- Preliminary project plan for migration

# Pre-Upgrade Checklist

This checklist can be used in addition to that found in the installation guide. It lists steps to prepare for upgrade. It is in order of earliest tasks to tasks done just before upgrading.

## Pre-Upgrade Checklist

- 1 Install the most recent version of the operating system, or at least the certified release for ASE.
- 2 Check that you have enough disk space. See your installation guide for requirements.
- 3 Check that each database has at least 10% free space. This is needed by the upgrade.
- 4 Check your applications to assess impact from SQL changes. See Chapter 5, “Implement: Making Required Application Changes”.
- 5 Check for new reserved words, including those in scripts and applications. See the section “New Reserved Words” in *Chapter 6* of this manual.
- 6 Replace obsolete environment variables in applications and scripts. See the installation guide or the white paper *Technical Issues in ASE 11.9.x and 12.0 Upgrade* at <http://www.sybase.com/detail/1,3693,1011767,00.html>.

- 7 Record size and device fragment information for all databases. You can query the `sysdevices` and `sysusages` tables for this information.
- 8 Make a note of the default character set and sort order. See the System Administration Guide at [http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/asesag/@Generic\\_\\_BookView?](http://manuals.sybase.com:80/onlinebooks/group-as/asg1200e/asesag/@Generic__BookView?) for information about character sets and sort order.
- 9 Verify that `sybssystemprocs` is large enough and that all device fragments, if there are multiple fragments, contain both log and data segments. See the section “`sybssystemdb`” in Chapter 6, “Implement: Making Database Administration Changes”.

Just before upgrading, be sure to do the following:

- 1 Set `master` as the default database for “sa”.
- 2 If the locations of your error log and configuration file are not the default locations, change them to the default location.
- 3 Back up the following:
  - All databases
  - NT Devices
  - Sybase directory
  - (NT) Sybase registry
- 4 `bcp out` critical system tables:
  - `Syslogins`
  - `Sysloginroles`
  - `Sysdatabases`
  - `Sysusages`
  - `Sysdevices`



- 5 On NT, if server is an automatic service, change the service type for ASE to manual.
- 6 Stop mirroring.
- 7 **Disable auditing:** sp\_auditoption "enable auditing", "off"
- 8 Set trunc log on chkpt **off** in all databases.
- 9 Save the server configuration (.cfg) file.

