



ArcSDE® Configuration and Tuning Guide for Microsoft® SQL Server™

ArcGIS™ 8.3

Copyright © 1986–2002 ESRI

All Rights Reserved.

Printed in the United States of America.

The information contained in this document is the exclusive property of ESRI. This work is protected under United States copyright law and the copyright laws of the given countries of origin and applicable international laws, treaties, and/or conventions. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying or recording, or by any information storage or retrieval system, except as expressly permitted in writing by ESRI. All requests should be sent to Attention: Contracts Manager, ESRI, 380 New York Street, Redlands, CA 92373, USA.

The information contained in this document is subject to change without notice.

RESTRICTED/LIMITED RIGHTS LEGEND

U.S. Government Restricted/Limited Rights: Any software, documentation, and/or data delivered hereunder is subject to the terms of the License Agreement. In no event shall the Government acquire greater than RESTRICTED/LIMITED RIGHTS. At a minimum, use, duplication, or disclosure by the Government is subject to restrictions as set forth in FAR §52.227-14 Alternates I, II, and III (JUN 1987); FAR §52.227-19 (JUN 1987); and/or FAR §12.211/12.212 [Commercial Technical Data/Computer Software]; DFARS §252.227-7015 (NOV 1995) [Technical Data]; and/or DFARS §227.7202 [Computer Software], as applicable. Contractor/Manufacturer is ESRI, 380 New York Street, Redlands, CA 92373-8100, USA.

ESRI, MapObjects, ArcView, ArcIMS, SDE, and the ESRI globe logo are trademarks of ESRI, registered in the United States and certain other countries; registration is pending in the European Community. ArcSDE, ArcInfo Librarian, Spatial Database Engine, ArcCatalog, ArcToolbox, ArcMap, ArcGIS, ArcStorm, ArcInfo, ArcObjects, ArcExplorer, ArcEditor, and the ArcInfo logo are trademarks and www.esri.com is a service mark of ESRI.

The names of other companies and products mentioned herein are trademarks or registered trademarks of their respective trademark owners.

CHAPTER 1

Getting started

Installing ArcSDE™ and loading data into a database are arguably simple processes, especially when using tools such as ArcGIS™ software's ArcCatalog™ or Microsoft SQL Server 2000's Enterprise Manager. So, why is there a configuration and tuning guide? While installation and data loading can be relatively straightforward, relational database management systems are dynamic. Relational databases offer highly efficient storage, searching, and modeling mechanisms that can serve many users simultaneously. It takes some effort to build and maintain a database that performs optimally. This book details how to manage your ArcSDE instance within this environment from setting up a two- or three-tiered server to serving your spatial data to multiple users and leveraging SQL Server advanced functionality.

Managing databases and logins

Microsoft SQL Server allows for multiple databases. Databases can be mobile, detached, and reattached to other servers using backup and restore or `sp_detach_db` and `sp_attach_db`. Databases can be replicated to other servers, can be used as standby servers, or can employ log shipping to distribute logs in backup environments.

Logins access databases as users. ArcSDE 8.3 has special requirements for managing users and databases. You can use integrated (Windows NT or Windows 2000) logins or SQL Server authenticated logins.

ArcSDE data is stored in SQL Server databases in tables, managed with stored procedures and triggers, and accessed with indexes. These entities are used in conjunction to maximize query performance and minimize database maintenance.

Arranging your data

Every table and index created in a database has to be stored somewhere. How you store these tables and indexes can affect how well your database performs under multiuser scenarios and complex queries.

The DBTUNE table

How is the placement of the tables and indexes controlled? ArcSDE uses storage parameters listed in the DBTUNE table that define how and where tables and indexes are created and stored. Storage parameters are grouped into configuration keywords that identify how and where a particular table or index is created.

Prior to ArcSDE 8.1, the configuration keywords were stored in the dbtune.sde file maintained under the ArcSDE etc directory. The dbtune.sde file is still used by ArcSDE 8.3 as the initial source. When the ArcSDE 8.3 sdesetupmssql.exe command executes, the configuration parameters are read from the dbtune.sde file and written to the DBTUNE table. You can import existing dbtune files into the ArcSDE dbtune table using the sdedbtune command. You can make direct edits to this table using the SQL Server Enterprise Manager.

ArcSDE 8.3 simplifies storage parameters by using actual t-sql CREATE TABLE and CREATE PROCEDURE storage parameters.

The ArcSDE 8.3 installation creates the DBTUNE table. If the dbtune.sde file is absent or empty, sdesetupmssql.exe creates the DBTUNE table and populates it with default configuration keywords representing the minimum ArcSDE configuration.

The ArcSDE 8.3 dbtune table has several default settings designed to provide a fast, scalable server “out of the box.” These settings influence clustered indexes, text in row, and network packet size.

In almost all cases, you will populate the table with specific storage parameters for your database. Chapter 3, ‘Configuring DBTUNE storage parameters’, describes the DBTUNE table and all possible ArcSDE storage parameters and default configuration keywords in detail.

Connecting to Microsoft SQL Server


The ArcSDE server provides two different ways to connect to your Microsoft SQL Server 2000 instance. ArcSDE clients may either connect to the ArcSDE service or directly to the database.

The traditional ArcSDE service connection method uses a three-tiered architecture. ArcSDE clients connect to the ArcSDE service, which then spawns a dedicated *gsrvr* process that connects to the database server. The *gsrvr* process brokers the spatial data between the ArcSDE client and the database server. The ArcSDE service and the *gsrvr* processes typically reside on the same machine as the database server, while ArcSDE clients typically operate on remote machines.

The direct connection method uses a two-tiered architecture. The ArcSDE clients connect directly to the Microsoft SQL Server. In this method, the ArcSDE *gsrvr* middle tier has been moved into the ArcSDE client. The ArcSDE clients still run on remote desktops, while the database server runs on a server class machine.

It is possible to use a combination of these two methods on any given database server. With a mixed configuration, some of the ArcSDE clients may be connected directly to a SQL Server, while others connect to an ArcSDE service.

Figure 1.1: With the traditional three-tiered system the ArcSDE client applications connect to the ArcSDE service, which in turn connects to the server. Within the two-tiered architecture the ArcSDE client applications connect directly to the server. Under a mixed architecture some of the ArcSDE client applications may connect directly to the ArcSDE service, while others connect directly to the server.



CHAPTER 2

Configuring SQL Server

Microsoft SQL Server is easy to set up and configure. SQL Server has outstanding management tools and wizards to assist you in performing database functions. Due to its ease of use and rich user interface, there are many possible configurations with SQL Server. This chapter guides you through ArcSDE-specific rules for creating and managing databases, logins, and transaction logs.

Introduction

This chapter picks up where a new ArcSDE installation leaves off. You've installed ArcSDE, let it create an sde database and login/user for you. Now you must consider how you should configure SQL Server for ArcSDE, design databases, manage security accounts, and understand the MS SQL Server transaction log.

To properly setup ArcSDE for MS SQL Server you must:

1. Make some configuration changes to MS SQL Server. These are detailed in the section entitled "Step 1: Installing and configuring Microsoft SQL Server."
2. Design your database or databases. You can put all your data into the 'sde' database or create new databases for your spatial data. This is covered under the section "Step 2: Creating and Managing databases and filegroups."
3. Design and configure your security. You must create logins and add them to the sde and other databases as users. You can use Windows Authentication as well. Section "Step 3: Managing Security" describes this.
4. Understand the MS SQL Server Transaction Log. The transaction log is a serial record of most activity against a database. The recovery process may employ it. Section "Step 4: Understanding the SQL Server Transaction Log" discusses this.

Step 1: Installing and Configuring Microsoft SQL Server

Installation and configuration of Microsoft SQL Server are simple processes. The installation of SQL Server requires one configuration change.

Installation of MS SQL Server:

1. **Use Mixed-Mode Security:** For the ArcSDE post-installation to work, you must choose mixed mode security. This specifies that your server supports both Windows NT/Windows 2000 logins and SQL Server Authenticated logins. ArcSDE 8.3 for Microsoft SQL Server supports both modes of security. If you want to operate in windows-only mode, you cannot run the ArcSDE post-installation utility. See directions below for windows-only installations. Security modes can be changed within SQL Server through the server properties in the Enterprise Manager. Right

click your server in the Enterprise Manager, select properties and then the security tab. Enable mixed mode security.

2. **Codepage:** Choose a codepage that is appropriate for your language. ArcSDE will work with most codepages, but is not presently supported running in mixed code page environments. This means that all your spatial databases should have the same collation as the 'sde' database.

MS SQL Server Serverwide Configuration

1. **Limit SQL Server Memory Consumption:** Do not allow MS SQL Server access to all your server's memory. By default, SQL Server will use as much memory as it requires. On busy systems, SQL Server can force other applications, including ArcSDE, to page, compromising their response time and scalability. Limit SQL Server's memory consumption by using a fixed memory size or setting a maximum memory size. To do this, right click the server in the Enterprise Manager and select properties. In the next form, select the memory tab. Hard code a memory limit or set an upper boundary to 50% of total ram. Adjust this number up until your server begins to swap and response time slows. On very busy systems, 50% may be too much.
2. **Use Lightweight Pooling:** On busy systems, configure SQL Server to use "lightweight pooling" or Windows NT Fibers. Fibers, which run as processes within a thread, reduce the number of thread context switches the processors must make. Access this setting in the Enterprise Manager's server properties under the Processor tab, or through sp_configure. See the MS SQL Server Books Online for more information.
3. **Boost SharedSection in the Registry:** This is especially true for ArcSDE servers supporting large ArcIMS sites. The setting controls the heap size for non-interactive desktop apps – the amount of memory available to spawn and support gsrvt (ArcSDE client connections). Locate this setting in the registry under: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\SubSystems under the "Windows" entry. This entry contains a string, contained in which is SharedSection=1024,3072,512. Boost the last entry to 1024. See the book "Managing ArcSDE Services" for more information.
4. **CPUs:** If your database will support many multi-versioned editors, consider that versioned queries are cpu intensive. Consider buying more cpu than less. Make sure you consult the ESRI System Integration system sizing tools before you buy a server to support this type of system.
5. **Number of SQL Server Licenses:** This depends upon which licensing mode you employ. If you choose client access licensing, you will need as many licenses as users who will connect simultaneously, plus one for the ArcSDE application server (the ArcSDE Service).

Support for SQL Server Named Instances

At MS SQL Server 2000, you can install multiple SQL Servers onto the same server. Multi-instance SQL Servers are known as "Named Instances." ArcSDE supports named instances in the same manner as a default instance.

To configure ArcSDE to work with a Named Instance, **do one of the following:**

1. During the ArcSDE Post-Install, type in the named instance whenever you are prompted for a datasource.
2. Use the sdeservice -o create tool, located in %sdehome%\bin to create a service pointing at a particular Named Instance of SQL Server. For example:

```
Sdeservice -o create -l <license server> -p <sde password> -s  
<named instance here> ...
```

See Appendix B for details on how to use the sdeservice command to identify a datasource.

Step 2: Creating and Managing databases and filegroups

Once you have installed the ArcSDE software, created an sde database and user, and made some SQL Server serverwide configuration changes, you will begin to design databases to support your applications. You can either use a single database (the 'sde' database) or several databases to store your data. You can employ filegroups and files to create a simple non-raid striping strategy, or employ a disk array for full data striping.

ArcSDE Multi-Database Model

If you choose to store your data in multiple databases, keep these rules in mind:

1. **Feature classes (layers) and rasters span database boundaries.** This means when a user connects to SQL Server via ArcSDE, that user is presented with a list of feature classes and rasters she has privilege to in **all** databases. **You can disable this behavior by setting the CROSS_DB_QUERY_FILTER dbtune parameter to 1 in the SDE_dbtune table.**
2. **Non-spatially enabled tables (registered or not) are filtered across database boundaries.** For example, if the cad user has two DBMS tables, income_avg in the sde database and traffic_stats in the county_hiways database, the cad user must explicitly connect to the county_hiways database to view the traffic_stats table.
3. DML (select, insert, update, delete) statements are allowed across databases. For example, you connect as user 'cad' and can select, insert, update, and delete data from any spatially enabled table (a layer or feature class) in any database.
4. DDL (create table, alter table, drop table) statements cannot occur across databases. If you connect as user 'cad' to the 'county_hiways' database, you can only load data (shp2sde, the ArcCatalog shapefile to geodatabase utility) into the 'county_hiways' database.
5. If you do not specify a database when you connect to ArcSDE, you connect to the sde database. The default database is sde, not a user's default database.
6. The 'sde' database is mandatory. It serves as the "master" database of the ArcSDE application.

What is CROSS_DB_QUERY_FILTER?

CROSS_DB_QUERY_FILTER is a new dbtune table parameter that allows an ArcSDE service to disable feature class and raster cross database query behavior. By default, it is set to 0 meaning allow cross database querying of feature classes and rasters. If you set it to 1, your connected clients can only work with feature classes and rasters in the presently connected database.

Why should you enable CROSS_DB_QUERY_FILTER?

1. To speed up your connection and layer list times.
2. To make your database a bit more secure – what they can't see can't hurt them.

Why should you disable CROSS_DB_QUERY_FILTER?

1. Cross database queries are useful – fewer reconnects to SQL Server to fetch data in another database.
2. Reduce the number of gsrvers spawned by an ArcIMS server. Each time you connect to a different database with ArcIMS, the spatial server will create new ArcSDE connections (new gsrvers). If you connect to the same database (even if you fetch data from a different database), the spatial server will reuse any existing gsrvers.

How Many Databases?

Now you must decide whether or not to employ a multi-database strategy or put all your spatial data into the 'sde' database.

Using only the 'sde' database – decision criteria

1. Administrative ease. If you have all your spatial data in one database, then your recovery and security strategies are less complex. Furthermore, you and your users will not have to contend with the ArcSDE multi-database model and its overhead. Finally, you'll have to backup only one database. If you have more than one database, you must always backup all your databases at the same time to keep the sde database in sync with the others.
2. Total number of tables will not be large. If you will have fewer than 100 total spatially enabled tables, put them all into the sde database. Remember, there are 3-6 tables per feature class, and 6 tables per individual raster (meaning if you store one image file as one raster layer). Fewer tables in a database means it will take less time to determine what the connecting user can access.

Using multiple databases – decision criteria

1. Very active read-write system. If you will serve data to more than a handful of users, employ multiple databases. Read-write ArcSDE servers will most likely use the sde_logfile and sde_logfile_data tables in the sde database and the versioning delta tables in non-sde databases. Spreading data into different databases (hence files) and hopefully across disks will lessen i/o contention. Servers serving data to larger scale, integrated ArcIMS-ArcGIS clients, or sites using the Metadata Server should consider using multiple databases.
2. Cleaner database design. It makes for a cleaner design to store data "thematically" into separate databases. For example, elevation data stored in the elevation database, roads, highways and railroads in the transportation database.
3. Scalable systems. If you have a large scale server with a lot of memory, disk and cpu power, administration of multiple databases becomes easier.
4. Raster and vector data. As the number of tables increases in a database, fetching performance against Raster layers will begin to decline for non-data owners because ArcSDE does not cache raster table permissions. There are three ways to store rasters in ArcSDE:
 - a. Mosaics: rasters that have the same cell registration can be mosaicked to one another creating a single large, seamless image. The advantage here is that the number of individual tables in the dbms is reduced. All ArcSDE raster clients support this format.
 - b. Multiple rasters in a single raster table: Using the sderaster tool, you can load multiple rasters of any type into a single raster layer. The images are not mosaicked but still loaded into the same set of raster dbms tables. Images within the table are identified by their image id (-v) value.
 - c. Multiple rasters in multiple dbms raster tables: This is equivalent to loading each raster image file into its own raster dbms table. As there are

5 dbms tables per raster layer, this solution will quickly create many tables in the dbms. This strategy should be avoided whenever possible in favor of points a and b.

Follow these general guidelines with regard to raster and vector when designing databases:

1. Create a separate database for your raster data if you will load many images or have many vector layers. Set `CROSS_DB_QUERY_FILTER` to 1 in the dbtune table. Keep in mind that you will most likely spawn extra gsrvers on the server.
2. If will have relatively few vector and raster layers, put them into the same database.
3. Mosaic or use the multiple rasters per raster layer whenever possible.

Client Side Considerations

You will quickly learn that what you do on the client very much matters to how your ArcSDE Server performs. Henceforth, your decision on whether to use multiple databases also depends upon what your clients will be doing.

1. **ArcIMS** – If ArcIMS servers see a different database in a connection string, then they will spawn a new set of gsrvers to support that map service. If you have a distributed ArcIMS solution with multiple ArcIMS servers hosting many virtual servers, you may exhaust the windows non-interactive desktop memory heap. You have three options:
 - a. Put all your data into a single database like the 'sde' database.
 - b. Always connect to the same database, but fetch your data from another database. This implies that you have not changed the `CROSS_DB_QUERY_FILTER` setting in the `SDE_dbtune` table.
 - c. Increase the `SharedSection` parameter listed above under MS SQL Server Serverwide Configuration.
2. **ArcGIS** – ArcGIS, by default, writes to an `sde_logfile` and `sde_logfile_data` table if more than 100 features are selected. The record of these writes ends up in the transaction log. Read-write ArcGIS clients will also hit the `sde_states`, `sde_tables_modified`, `sde_lineages_modified`, and the `sde_lineages` tables in addition to the multi-versioned tables that comprise a feature class. You should:
 - a. Use multiple databases. ArcGIS supports the multi-database model very well. You can create RelationshipClasses across database boundaries, for example, only if you do not disable `CROSS_DB_QUERY_FILTER`.
 - b. Set a high `SelectionThreshold` in your registry to store more selected features locally instead of on the sde server. Setting this to 1000 features is reasonable for many client systems. Keep in mind that this setting requires more processing on the client side, so make sure you have enough memory and cpu to handle it.
3. **Integrated Systems** – Metadata Server. The Metadata server uses the `sde_logfiles` and `sde_logfile_data` tables extensively. Therefore, if your server is in an integrated ArcGIS-ArcIMS environment with the Metadata server, you should use multiple databases and set a high `SelectionThreshold`.

Guidelines for Database Creation

Now that you have decided how many databases to create, here are some guidelines for creating any database to store spatial data.

1. If the Post-Installer created your sde database, you should increase the autogrowth increment of both the database and transaction log. While you're at it, make a backup of it.
2. Store all your data files, transaction logs and tempdb away from the paging file, unless you're sure your server will never page.
3. Separate your data files from your transaction logs.
4. If you have a RAID 5 solution, keep your transaction logs and entire tempdb out of the disk array.
5. Autogrow the data files in large increments.
6. Leave `AUTO_CREATE_STATISTICS` and `AUTO_UPDATE_STATISTICS` enabled. Disable auto shrink.
7. After you create a database, make a backup of it to force SQL Server to manage the transaction log.
8. Employ a hardware striping solution, favoring RAID (disk array) over Filegroups and Files.
9. Employ data segregation strategies (keeping tables from indexes, or certain types of tables from other tables) only if you are certain it will improve performance or alleviate administrative burdens.

Hardware Striping Solutions

For production ArcSDE systems serving many simultaneous users and storing large amounts of data, you must employ a hardware striping solution. Your best disk/data organization strategies involve spreading your data across multiple disks. With data spread across multiple disks, more spindles actively search for it, increasing your fetch times. Spreading your data across many disks will decrease disk contention as well. However, there is a saturation point in which too many disks will slow down a query. There are two main ways of achieving striping: Filegroups and RAID. You can also combine the two – create filegroups within disk arrays.

Filegroups and Files

Filegroups and files compose a database. Data resides in files that are grouped by filegroups. The database's primary file of the primary filegroup has a *.mdf extension, while secondary files use *.ndf. Log files are stored in *.ldf files. A filegroups group files into administrative units. Database files cannot span disks, filegroups can.

Files fill proportionately according to their membership in a filegroup. For example, if three files comprise a filegroup, and file A is 20 Mb, file B is 40 Mb, and file C is 60 Mb, one extent is allocated from file A, 2 from B, and 3 from C whenever space is requested from a filegroup. All database objects are created by default on the primary filegroup. Placement onto other, non-default, filegroups is controlled with the "on" keyword of the "create" statement. ArcSDE controls placement through the use of the `SDE_dbtune` table in the same format – by adding the keyword "on" to your storage parameters.

Not all ArcSDE installations will be on multiprocessor server class machines with disk arrays or many disks. ArcSDE is certified to scale from Microsoft Data Engine (MSDE)

to SQL Server Enterprise Edition. If you are serving data to a few clients on a single processor machine with one or two disk drives, employing filegroups offers no advantage.

RAID

Redundant Arrays of Inexpensive or Independent Disks (RAID) boosts fetching performance by striping data into slices across multiple disks in the disk array. By spreading data across multiple disks, all disks share the burden of I/O operations thus reducing the chance of a bottleneck occurring on one disk. RAID's performance increases as you add disks to the array. The operating system and database will see only one volume, a logical representation of the entire disk array.

Striping Design Considerations

There are two main striping designs, *smear* and *plaid*.

1. Smear, in the SQL Server context, signifies the creation of a single filegroup with one large file in a disk array, or a single filegroup with multiple files spread across several disks. In this configuration, the RAID solution will most likely perform best. When you load your data, you lump all your data and indexes into the same filegroup and leave it to SQL Server to fetch it. There is no segregation of tables and indexes, nor hot and cold tables. Smear's advantages include ease of administration and possibly fewer disks to purchase.
2. Plaid connotes the strategic and calculated placement of tables and indexes onto disk arrays or filegroups that are independent of each other. The segregation of components tends to reduce the chance for i/o contention. In such a configuration, you create different filegroups for your indexes and tables, and perhaps types of tables (business, feature, and spatial, for example). In a RAID solution, you would have different disk arrays for your tables and indexes, or even different filegroups within your disk arrays. Plaid's advantages are that on very large databases with many concurrent users, data segregation may enhance scalability.

There is no concrete performance data indicating that on small to mid-size databases one solution performs better than another. For very large databases, Plaid may be the best since more data must be fetched from disk. The clustered index is used overwhelmingly in all types of fetches. Many operations do not use the non-clustered indexes at all.

One Final Word

You'll rarely have to worry about fetching from disk if you have enough ram to cache the database or if the database is small enough to be cached.

Example Databases

This section offers some example spatial database design strategies based upon client usage considerations and ease of administration. Cost of components is not discussed.

Supporting primarily ArcIMS clients with both read-only vector and raster data

1. If you have less than 30 –40 combined vector and raster layers:
 - a. Create the sde database on a disk array. Autogrow your datafiles and transaction log in large increments.
 - b. Load the raster data into mosaics or the “multiple rasters per raster layer format” (reducing the number of tables in the database).
2. If you have more than 40 combined vector and raster layers:
 - a. Create the sde database on a disk array. Place the transaction log onto a different disk.
 - b. Create a Raster database on a disk array. Set both initial size and autogrowth parameters to large increments.
 - c. Create a database or series of databases for your vector data, organizing them by project or thematically. Try to keep the total number of databases low. Place the transaction log onto a different disk and autogrow all datafiles and transaction log files in large increments.
3. Set CROSS_DB_QUERY_FILTER to 1 in SDE_Db tune. If you have no raster data, skip this section.
4. Calculate the number of gsrvs that will run on the ArcSDE server as:
 - a. $2(\text{number of spatial servers}) + 2(\text{number of query servers})$
 - b. multiply this number by 2 for each unique database listed in your <SDEWORKSPACE> tags in your map service files.

The presumes that you have not changed the default ArcIMS thread settings. By default, ArcIMS runs two instances of the spatial server per installed spatial server, and two instances of the query server per installed query server. For example, you have one spatial server and one query server configured, and you author one image service. The number of gsrvs created on the server will be 4 – 1 per instance of the spatial server (2 gsrvs), and 1 per instance of the query server (2 more gsrvs). If you author another axl that connects to another database, you’ll get an additional 2 gsrvs from the spatial server.

If you begin creating a lot of gsrvs on the ArcSDE server, you may exhaust heap memory, forcing you to increase the windows SharedSection setting listed in Step 1 above.

Supporting primarily read-only ArcGIS clients with both vector and raster data

1. Create the sde database within a disk array. Set autogrow to a large increment. Create or move the sde transaction log on another disk outside the disk array.
2. Backup the sde database.
3. Create a raster database on the disk array. Set it to a large initial size with a large autogrowth rate commensurate with the size of your data.
4. Create a database or series of databases for your vector data, organizing them by project or thematically. Set each database to a large initial size and growth increment commensurate with the size of your data.

5. Load your rasters into as few mosaicked seamless images as possible.
6. Set CROSS_DB_QUERY_FILTER to 1 in the SDE_Dbtime table.
7. If your workflow includes large selection sets (routinely greater than 5000), customize the SDE_Dbtime:
 - a. Add entries for logfile parameters under the LOGFILE_DEFAULTS keyword. See the Chapter 3 for details. Increase all FILLFACTOR configuration strings to 95 percent.
 - b. If you have a lot of hardware, a very busy system, and a very large database, consider placing the logfile tables and indexes onto a separate array.
8. Add the SelectionThreshold setting to all ArcGIS client machines. For a workflow including large selection sets, boost the setting to 5000 machine permitting.

Supporting read-write ArcGIS clients

1. Create the sde database within a disk array. Set autogrow to a large increment. Create or move the sde transaction log on another disk outside the disk array.
2. Backup the sde database.
3. If you will support large selection sets (greater than 5000), add the logfile parameters and the LOGFILE_DEFAULTS keyword to SDE_Dbtime table. Increase FILLFACTOR config_strings to 95 percent.
 - a. If you have a lot of hardware, a very busy system, and a very large database, consider placing the logfile tables and indexes onto a separate array.
4. Create a database or series of databases for your vector data, organizing them by project or thematically. Set each database to a large initial size and growth increment commensurate with the size of your data.
5. Set CROSS_DB_QUERY_FILTER to 1 in the SDE_Dbtime table.
6. Add the SelectionThreshold setting to all ArcGIS client machines. For a workflow including large selection sets, boost the setting to 5000, machine permitting.

Supporting Large, Integrated Systems (ArcIMS, ArcGIS, MapObjects Java)

This example presumes your ArcSDE server will support ArcIMS map services, ArcIMS metadata explorer services, ArcGIS editors, and other clients like the ArcSDE java api or MapObjects.

1. Create the sde database within a disk array. Set large autogrowth increments. Move the sde transaction log to another disk, if possible, outside the disk array. Backup the sde database.
2. Add the logfile parameters and the LOGFILE_DEFAULTS keyword to SDE_Dbtime table. Increase FILLFACTOR config_strings to 95 percent.
 - a. Consider placing the logfile tables and indexes onto a separate array.
3. Customize the ims* SDE_Dbtime keywords to place ims_metadata objects outside the sde database.

7. Create a database or series of databases for your vector data, organizing them by project or thematically. Set each database to a large initial size and growth increment commensurate with the size of your data.
8. Set CROSS_DB_QUERY_FILTER to 1 in the SDE_Dbtime table.
9. Add the SelectionThreshold setting to all ArcGIS client machines. For a workflow including large selection sets, boost the setting to 5000, machine permitting.

Example: Using SQL to create an sde database

The easiest way to create a database, filegroups, and data files is to use the SQL Server Enterprise Manager. Expand your server, right-click the databases node, and select “New Database.” However, you can employ SQL scripts to automate this activity.

This example creates an sde database with one filegroup and one large datafile within a RAID array. The log file is created on a separate disk. The filegroups are created so that the ArcSDE system tables are left on the primary filegroup.

```
Create Database sde
On Primary                                /*created on primary filegroup */
( NAME = sde_dat,
  FILENAME = 'g:\data\sde.mdf',
  SIZE = 40,
  MAXSIZE = UNLIMITED,
  FILEGROWTH = 15% ),
FILEGROUP SpatialData                    /*specify another filegroup to
separate*/
( NAME = SpatialData_dat,                /*spatial data*/
  FILENAME = 'g:\data\SpatialData.ndf',
  SIZE = 400,
  MAXSIZE = UNLIMITED,
  FILEGROWTH = 400 )
LOG ON
( NAME = 'SDE_LOG',                      /*place log on different drive*/
  FILENAME = 'F:\logs\sde_log.ldf',
  SIZE = 200,
  MAXSIZE = UNLIMITED,
  FILEGROWTH = 200 )
GO

/*add logins to the server*/
sp_addlogin 'sde', 'esri.sde', 'sde'
go
sp_addlogin 'bob', 'esri.bob', 'sde'
go
sp_addlogin 'carto', 'world.round', 'sde'
go

/*switch to the sde database and
add users */
use sde
go
sp_grantdbaccess 'bob'
go
sp_grantdbaccess 'sde'
go
sp_grantdbaccess 'carto'
go

/*permissions*/
grant create table to sde
grant create procedure to sde
grant create view to sde

grant create table to bob
grant create table to carto
```

Step 3: Managing Security

Once you have designed and created your databases, you now must configure your security accounts. Create logins and users with specific functionality in mind, then assign them to roles and finally grant permission to data to those roles.

Introduction to SQL Server Security

Microsoft SQL Server has ‘logins’ and ‘users.’ A login is a named account that has access rights to the database server but not necessarily access to a database. A user is a login that has been granted access to a database.

There are two types of authentication: SQL Server authentication and Windows NT or Windows 2000 authentication. SQL Server authentication specifies that the named accounts are valid only in SQL Server. Windows NT or Windows 2000 authentication signifies that an operating system account is also valid within the database server.

MS SQL Server also supports roles. Roles are special groupings of users to which a specific privilege is granted. There are three types of roles: fixed server, database, and user-defined. Fixed server roles span the server and apply to all databases. Database and user-defined roles pertain to a specific database.

Logins and users: The rules for ArcSDE

Before you create logins and users for ArcSDE, you must understand the requirements for permissions and database access. Because you can have multiple databases in MS SQL Server and have different types of authentication, you should ensure that your logins adhere to these rules to prevent unexpected client behavior when accessing your spatial data. Remember that a database user is a login that has been granted access to that database.

1. The ‘sde’ user must have create table, create view, and create procedure within the ‘sde’ database. The sde user must be able to create the metadata tables and stored procedures that manage an ArcSDE instance. If you are using SQL Server 2000, the sde user must also have “create function” permission.
2. The ‘sde’ login must be granted access to any other database in which you will store spatial data.
3. You must add all logins that will access spatial data as users to the sde database and grant them create table permission. These users must be able to write their logfiles into the sde database and select against the sde metadata. NOTE: The first time a user accesses spatial data managed by ArcSDE, they create two logfile tables in the sde database. As long as these tables are not deleted, it is permissible to revoke their create table permission because they have created their logfile tables.
4. Any user that will own (create) data in any database through the ArcSDE application must have create table and create procedure in that database. They must have create table to create the tables comprising an ArcSDE feature class and create procedure to create objectid generation stored procedures that accompany a feature class.
5. You must explicitly connect to a database to create data in it, drop a table from it, or alter a table within it. In other words, DDL statements can only be issued on data if you have explicitly connected to the database storing that data. If you don’t specify a database name in your connection string, you connect, by default, to the sde database. We ignore a login’s default database!

6. You can connect to any database and edit data to which you have permissions in any other database.
7. You should have a separate account for all your application users to reduce contention against the logfile tables in the sde database. The logfile tables are used to store selection ID's in programs such as ArcGIS or ArcView® DBAccess. If you have multiple users logging into ArcSDE as the same account and then selecting features in their client programs, they may see unpredictable results or experience performance problems due to their sharing of a single logfile table.
8. You cannot use a SQL Server reserved word for a login or user name. Examples of reserved words are Max, National, and Count. See the SQL Server books online (BOL) for a complete listing.

ArcSDE users rule matrix

User	In SDE database	In other database
SDE User	create table, create view, create procedure, create function (SQL Server 2000)	must be added as a user
Other User	create table, create procedure if will own data	add as a user if will access data, create table and procedure if will own data

Designing your Security

Setting up your security follows this general process:

1. Creation of logins and assignment of those logins to databases. The logins become users.
2. Create database roles or use existing roles and add users to them
3. Grant permission to data to those roles.

Security Considerations

MS SQL Server security can be confusing. ArcSDE complicates it by requiring specific permissions. Here are some suggestions to simplify the process:

1. For simplest administration
 - a. Use SQL Server authentication only.
 - b. Do not add the sde login to any specific roles. The sde user must have create table, procedure, view and function (SQL2k only) in the sde database and have access to all other spatial databases.
 - c. Create individual sql server logins and add them to databases. If they need read permission to all data in a specific database, add them to the db_datareader role.
2. For tightest security:
 - a. Configure ArcSDE and SQL Server to use only windows logins
 - b. Set CROSS_DB_QUERY_FILTER to 1 in SDE_Db tune table
 - c. Grant no permissions on data to the "Public" role.

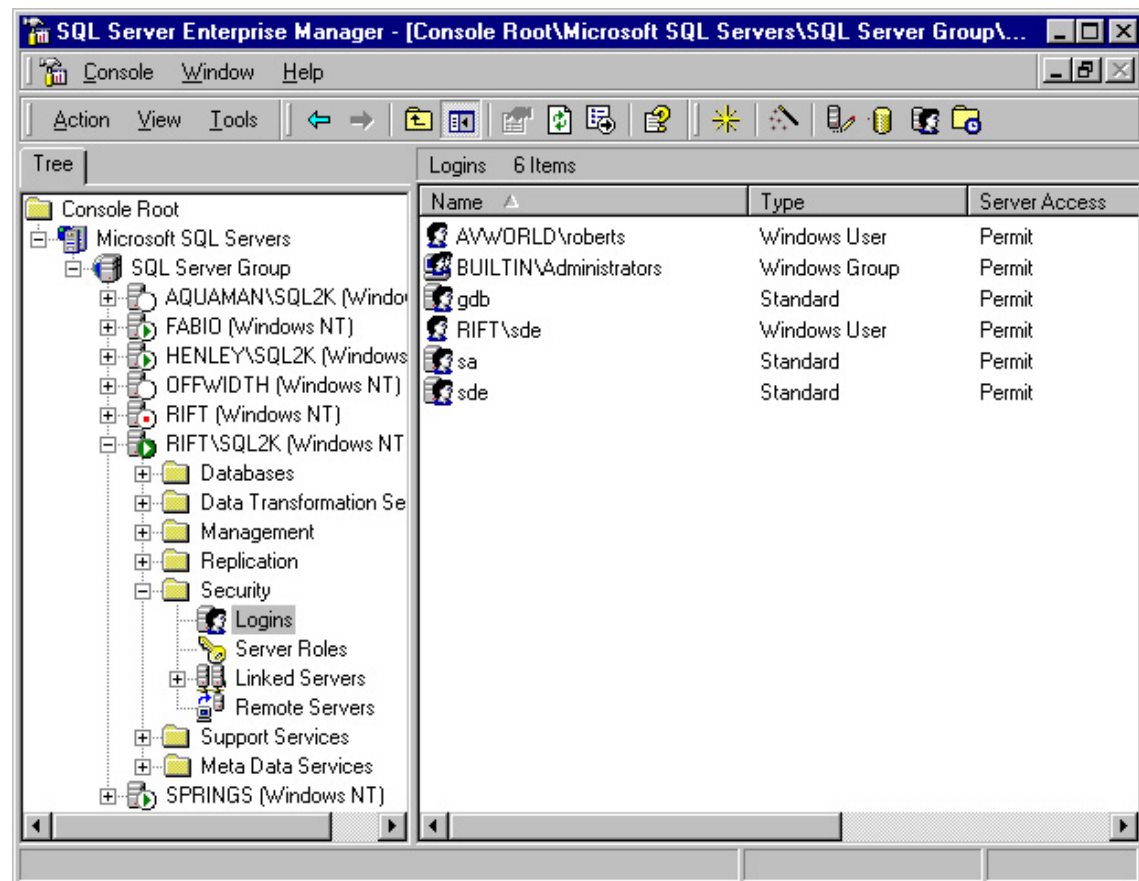
- d. Do not add any login to the db_owner database role.
 - e. Never save your username and password in an ArcGIS connection file.
 - f. Use direct connect for ArcIMS connections and input a junk string into the password field.
3. Avoid any use of fixed server roles unless you are very comfortable with them.
 4. Preserve the 'sde' login as the system administrator; do not use it for other operations. Do not add the sde login to the sysadmin fixed server role.

1. Creating a new SQL Server authenticated login for use in ArcSDE

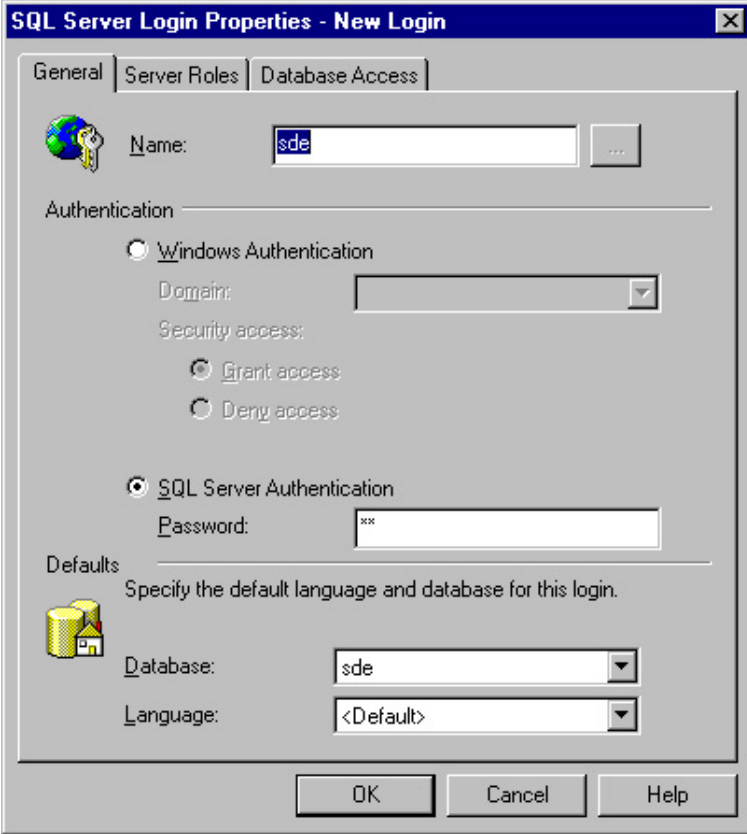
Keep in mind there are multiple ways to do this. SQL Server 7 and 2000 also differ slightly, but the idea is the same: you must first create or add a new login to the relational database management system (RDBMS) before you can access data as that login. Each topic listed below describes the entire process from creating the login to granting it access to a database or databases, and granting that user privileges within a database.

Enterprise Manager

1. Open the SQL Server Enterprise Manager, expand your server's branch in the treeview, expand the "security" node, and select the logins topic.




2. In the details panel on the right, right-click and select "new login."



SQL Server Login Properties - New Login

General | Server Roles | Database Access

 Name:

Authentication

☐ Windows Authentication

Domain:

Security access:


☐ Grant access

☐ Deny access

☒ SQL Server Authentication

Password:

Defaults

 Specify the default language and database for this login.

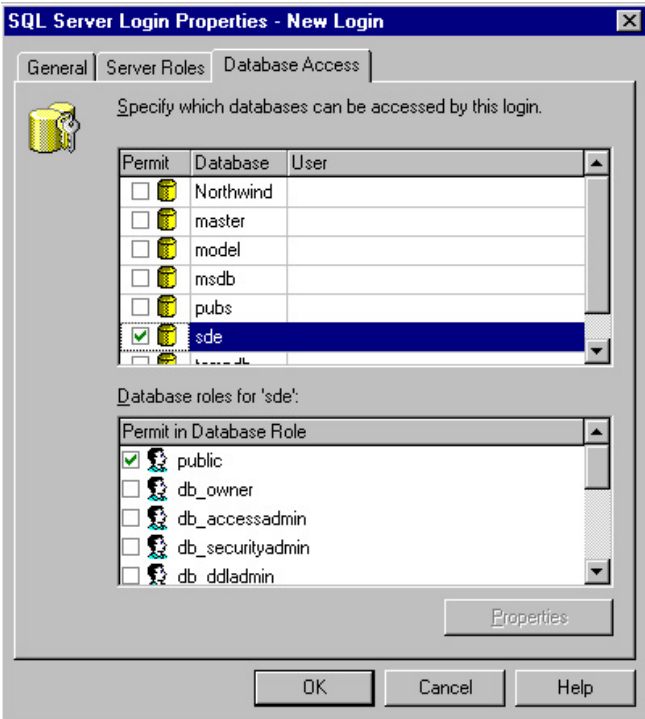
Database:

Language:

OK Cancel Help


Select the SQL Server Authentication radio button, name the login, and provide a password. Assign a default database.

- Now click the Database Access tab. Scroll through the list of databases and select the sde database and any other database you want to add this login to as a user.



SQL Server Login Properties - New Login

General | Server Roles | Database Access

 Specify which databases can be accessed by this login.

Permit	Database	User
<input type="checkbox"/>	Northwind	
<input type="checkbox"/>	master	
<input type="checkbox"/>	model	
<input type="checkbox"/>	msdb	
<input type="checkbox"/>	pubs	
<input checked="" type="checkbox"/>	sde	
<input type="checkbox"/>	tempdb	

Database roles for 'sde':

Permit in Database Role	Role
<input checked="" type="checkbox"/>	public
<input type="checkbox"/>	db_owner
<input type="checkbox"/>	db_accessadmin
<input type="checkbox"/>	db_securityadmin
<input type="checkbox"/>	db_ddladmin

Properties

OK Cancel Help

4. Click OK and you'll be prompted to confirm your password. The process both creates the new login and grants it access to a database or databases (adds that login as a user to a database).
5. Now expand the databases tree in the treeview and right-click the sde database (or any other database you've added that login to as a user). Select Properties from the popup menu. In the properties form, select the Permissions tab. You'll see a grid of users of this database with columns listing permissions. Check off "create table" and "create procedure" for a user who will own data in this database. Check off "create table" for any user in the sde database. Click OK.

Transact SQL

Scripting the creation of your logins and users is much faster than using the Enterprise Manager. The same rules apply as with the Enterprise Manager. Below is a SQL script that will demonstrate this idea.

```
/*
        Add sde and non-sde login and user
*/

/*
        Step1: Add the logins to the server using
        the sp_addlogin stored procedure.
        sp_addlogin '<login_name>','<pass>','<default_db>'
*/

use master
go
sp_addlogin 'sde','spatial.data','sde'
go
sp_addlogin 'arthur','dent','sde'
go

/*
        Step2: Grant the login access to the sde database
        using sp_grantdbaccess. Grant the user create table
        and stored procedure. You must switch to the sde
        database to do this.
*/

use sde
go
sp_grantdbaccess 'sde'
go
sp_grantdbaccess 'arthur'
go
grant create table to sde
grant create table to arthur
grant create procedure to sde
grant create view to sde
grant create function to sde --SQL Server 2k only!
grant create procedure to arthur
go
```

Using Windows 2000 or Windows NT logins with ArcSDE

ArcSDE 8.3 for MS SQL Server supports both SQL Server and Windows authentication.

Windows users are parts of domains and groups. You can add individual windows users SQL Server and connect with them to ArcSDE data stores. There are issues with this implementation as login names and user names can differ.

To use integrated logins with ArcSDE you must:

1. Add the domain user to a group on the ArcSDE server.
2. Add the domain user to the SQL Server as a new login.
3. Grant the new login access to a database.
4. Grant the login appropriate permissions within the database.

You don't have to perform points 2-4 if your Windows NT user is added to the machine and is a member of the Windows NT administration group on that machine.

NOTE: Windows NT groups are not supported.

Using Windows 2000 or NT 'sde' login

You can use an integrated 'sde' login instead of a SQL Server authenticated sde login. Keep in mind that the installer will always create a SQL Server 'sde' authenticated login. However, you don't have to run the postinstaller to set up your sde database.

To use an integrated sde login:

1. Add the domainname\sde login to the SQL Server logins collection.
2. Add the sde login to the sde database (and any other database you intend to use for spatial data storage) as a user.
3. Grant the sde login create table, view, stored procedure, and function (create function is necessary in SQL Server 2000 only).
4. Login to Windows as domainname\sde. Open a DOS prompt and switch directories to %SDEHOME%\bin. Run sdesetupmssql -o install as the sde user. Do not input a username or password.
sdesetupmssql -o install -s rift -H C:\sde81\sqlxe
5. Or, open a DOS prompt and switch directories to %SDEHOME%\bin. Run sdesetupmssql -o install as the sde user. Input a null username or password if your windows account has admin privileges or use the "sa" account.
sdesetupmssql -o install -s rift -H C:\sde81\sqlxe
sdesetupmssql -o install -s rift -H C:\sde81\sqlxe -u sa -p <sa password.
6. Run sdeservice -o create to create an sde service. Specify a null string for your password -p "" argument. Use the minimum options: -s, -l, -d, -H.
sdeservice -o create -s rift\sql2k -l Fabio -d SQLSERVER -H C:\sde81\sqlxe -i esri_sde -p ""
7. Start the sdeservice: net start esri_sde, alternatively, make a direct connection to it.

Connecting to an ArcSDE service as a Windows NT login

You can connect to an ArcSDE service without using a username or password with some ArcSDE clients. The examples below use the administration commands to demonstrate this functionality.

Sdelayer:

```
C:\>sdelayer -o describe -i sql82k
ArcSDE 8.1 Build 633 Tue Nov 21 22:30:10 PST 2000
Layer Administration Utility
```

```
-----
Database : VTEST
Table Owner : VTEST
Table Name : SANDIEGO_NET_JUNCTIONS
```

```
Spatial Column : SHAPE
Layer id : 8
Entities : npc
Layer Type : SDE
I/O Mode : NORMAL
User Privileges : SELECT, UPDATE, INSERT, DELETE
Layer Configuration: DEFAULTS
```

sdetable:

```
C:\>sdetable -o create -t authentication_test -d "colA string" -i sql82k
```

```
ArcSDE 8.1 Build 633 Tue Nov 21 22:30:10 PST 2000
Attribute Administration Utility
```

```
-----
Successfully created table authentication_test.
```

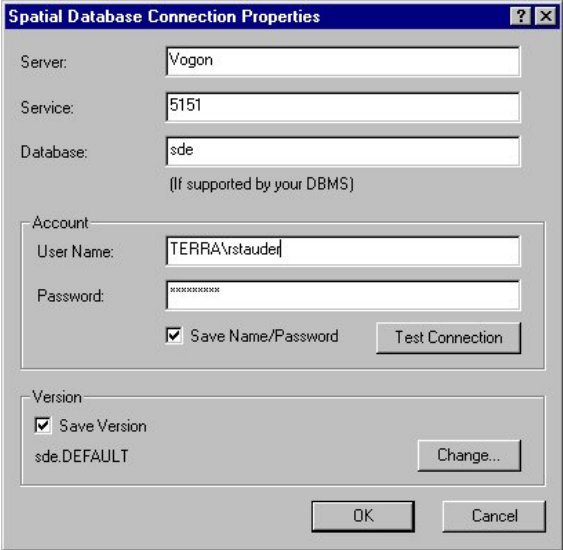
Sdeimport:

```
E:\world_sdex>sdeimport -o create -l test,shp -i sql82k -f cntry94
```

```
ArcSDE 8.1 Build 633 Tue Nov 21 22:30:10 PST 2000
SDEX File Import Administration Utility
```

```
-----
Importing SDEX from cntry94 ...
165 features read.
165 features stored.
```

ArcCatalog:



NOTE: Using integrated logins (Windows 2000 and NT) in direct connect

If you make a direct connection to ArcSDE using a Windows 2000 or NT user, your connection login and password will be taken from your operating system login.

2. Using roles

Once you have created logins and granted them user access to databases, you should create roles and assign those users to roles. Add your users to roles based upon operations they'll perform. For example:

- € Viewers: users that have read-only access to data.
- € Editors: users that will edit data in a database.
- € Developers: users that will create and deploy applications based upon data in a database.

There are four types of roles, database, user defined, fixed server, and application. Application roles are not supported. Database roles are pre-existing, sql server defined roles. User defined roles are created by an administrator within a database.

Adding a database user to a role simplifies administration by reducing the number of things you have to do per user. If you define a role that has create table and procedure and then add all your ArcSDE users to it, you have reduced the number of times you have to grant a privilege to each one.

Create a User Defined role with `sp_addrole` or the Enterprise Manager. For example:

```
sp_addrole 'Editors'
```

Will create a user defined role 'Editors' owned by the dbo user. You would now add users to this role with the MS SQL Server Enterprise Manager or the `sp_addrolemember` stored procedure:

```
sp_addrolemember 'Editors','bob'
```

Using database roles

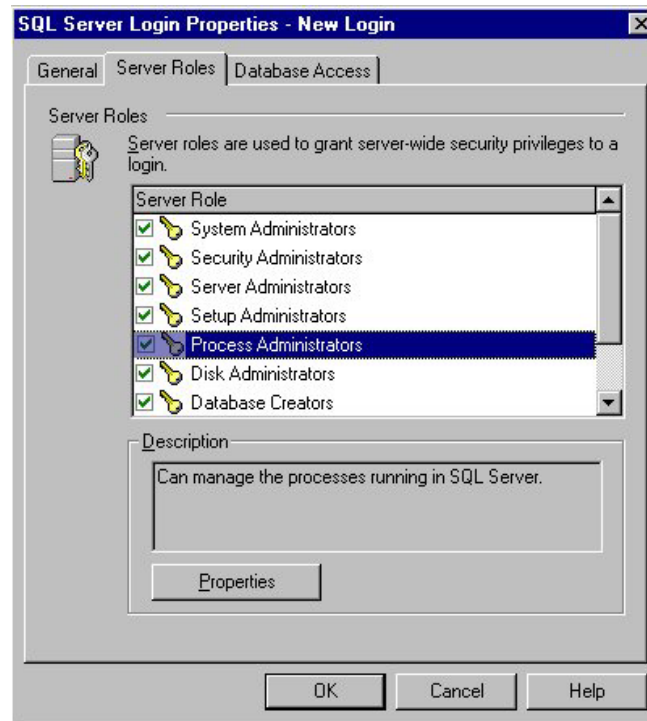
Database roles are pre-existing roles that carry implicit privileges and permissions. Public is a database role that all users belong to. Exercise extreme caution when using database roles as they can produce unexpected side effects. The advantage to using database roles is that they carry implicit permissions and privileges meaning less work for you to do.

When to use database roles:

- € Db_datareader: To grant select on all tables in a database to a role member.
- € Db_ddladmin: Permits a user to run any ddl statement in a database.
- € Db_owner: Make a user sysadmin of a database without changing the user name to "dbo."

Fixed server roles

Fixed server roles pertain to the server, and hence, all databases. There are two main things to consider here: if you add a user to the sysadmin fixed server role, anything that user creates in the database is owned by 'dbo' in any database, and you can add a user to multiple, sometimes conflicting, roles. This last point will assuredly cause problems when you connect and attempt to view (or load) your data. It is easy to do this, so make sure you don't. Figure 2.5 is a screenshot of SQL Server 2000 that shows this. Be careful you don't do this!

Figure 2.3: Adding a login to all fixed server roles—*Don't do this!*

3. Grant permission to data to those roles

Once you've created logins, added them as users to databases, and added those users to roles, it is either time to load data, or grant permission on that data to the roles. Chapter 4 covers data loading but before you can load and grant permissions, you must decide who should own the data, and henceforth do the granting.

Who should own the data?

Designate one user as the owner of all your data. Load all your data as that user and then grant permissions to that data to roles. The advantage of this system is that your database will not get cluttered with tables belonging to several different users.

If you load all your data and build ArcGIS complex features, and then add your dataowner to the sysadmin fixed server role, you may cause problems maintaining those complex features.

Granting Permission to Simple Features

Use the `sdelayer -o grant` or `-o revoke` command to grant or revoke a privilege (SELECT, INSERT, UPDATE, DELETE) to a role. For example, you have a role called 'Editors' in the sde database. To grant this role SELECT, INSERT, UPDATE, DELETE permissions against an ArcSDE feature class, 'county_hiways' -

```
sdelayer -o grant -U highways_editors -A SELECT, INSERT, UPDATE, DELETE -l
county_hiways,shape -u county -p big.bob -i esri_sde -s vagon
```

Granting Permission to Complex Features

As complex features are managed by the ArcGIS application, you must use it to grant permissions to such data to roles or individual users. Right click the feature class or

feature dataset in ArcCatalog and select Privileges. **NOTE:** A user or role must have insert, update and delete privilege to make multi-versioned edits.

Managing transaction logs

Transaction logs are composed of a physical log, virtual logs, and a logical log. The physical log represents the transaction log file or files on disk, while virtual logs are internal segments of those files. The logical log contains the active portion of the transaction log.

Transaction logs record transactions, data modifications, extent allocations and deallocations, and creation/destruction of tables and indexes. This record allows you to:

1. ROLLBACK changes made within a transaction.
2. During recovery, roll back incomplete transactions initiated before a database failure.
3. During recovery, roll forward modifications listed in the log but not yet written to disk.

Each record written to the log has a Log Sequence Number (LSN). All new records are written to the end of the log and have a higher LSN than its predecessor. The logical log will grow until it consumes all open space in the physical log. This will cause the physical log to autogrow if allowed or throw an error. If the database uses the simple recovery model, the logical log will truncate itself when automatic checkpoints occur. If the database uses the full or bulk-logged recovery model, log truncation will occur after a BACKUP LOG statement completes.

The logical log cannot truncate itself past the Minimum Log Sequence Number (MinLSN), usually the oldest uncommitted transaction recorded.

Transaction logs are used to recover a database in the full or bulk-logged recovery models. The simple recovery model only restores a database to its last backup. If you require point-in-time recovery of your data, use the full recovery model for your database. If your data is unchanging, use the simple recovery model. The bulk-logged recovery model performs minimal logging of certain operations and does not support point in time recovery.

If you have never made a backup of your databases, you'll notice that your transaction logs have grown to considerable size. ArcSDE Transaction logs grow due to data loading, editing, and making selections into sde_logfiles. When you make selections with ArcGIS clients such as ArcMap™, all selections exceeding 100 records are written into a database logfile. The continuous inserting and deleting from this table will make the database transaction log grow.

Backing up your database and logfile should control the size of the log by truncating it. If you have not made regular backups of the transaction log, you'll need to take action to reduce its size. To reduce the amount of space used within the physical logfile by the logical log, truncate it. To reduce the overall size of the physical logfile, shrink it.

Truncating transaction logs

Truncating a transaction log refers to the process of removing records from the log that are no longer necessary for database recovery. These LSNs of these records would precede the MinLSN of the active portion of the logical log. Truncating the log will not shrink the size of the physical log.

To truncate your log do one of the following:

1. Backup the database using either the Enterprise Manager or the Backup statement.

2. Backup the logfile using the Enterprise Manager or the Backup statement.
3. Use the simple recovery model. This model will automatically truncate your database whenever a checkpoint occurs or when the log fills to 70%.

Once you've issued the backup statement, all records preceding the MinLSN are marked for reuse, and the reported size of the transaction log's logical log is smaller. The following figures demonstrate this.

Figure 2.4: A transaction log after registering some very large feature classes with the Geodatabase





Figure 2.5: The same log after a database backup. The log has been truncated.



Shrinking the transaction log


Truncating the log marks all nonactive transactions as reusable but does not reclaim any space, only delaying any more growth to the physical log file. In the case shown above in Figure 2.5, the physical log file still consumes 315.8 MB even though only 11.4 MB are active. You may find that you need some of this extra space. Shrinking the transaction log will return allocated space to the operating system.

To shrink the transaction log do one of the following:

1. Use the Enterprise Manager's Shrink database utility: Right-click your database and select all tasks – Shrink Database. Input your desired resulting free space and click OK.
2. Use DBCC SHRINKDATABASE.

In the figure below, we used the Enterprise Manager to shrink the logfile to 49.7 MB.

Figure 2.6: The logfile's new size after shrinking it



Help! It won't shrink!

Under certain circumstances, you will not be able to shrink the size of the logfile. As the logical log progresses through serial lists of Log Sequence Numbers, the active portion of the log will arrive at the end of the physical file. You cannot return space preceding the MinLSN to the operating system. What you must do is advance the MinLSN to the beginning of the physical file by following this process:

1. Backup the database. This will mark the preceding records to the minimum log sequence as reusable. Check the size of the log file.
2. Execute some updates using t-sql. This should move the active portion of the log toward the beginning of the physical file.
3. Execute `dbcc shrinkdatabase (dbname)` to create a "shrinkpoint."
4. Backup the log file.

At this point you should see more free space in the logfile. Repeat a few times if necessary.

CHAPTER 3

Configuring DBTUNE storage parameters

Each ArcSDE service contains a metadata table, `SDE_dbtune`, that contains customization parameters for data loading and access. `SDE_dbtune` allows a database administrator to control data placement onto filegroups, index fillfactor, binary data inlining, and the filtering of feature classes and rasters across database boundaries. This chapter explains the purpose of `SDE_dbtune` and the keywords, parameters and configuration settings it contains.

The `SDE_dbtune` table

The `SDE_dbtune` table is created during setup of the ArcSDE software by the post-installation setup or the `sdesetupmssql.exe` command. The `SDE_dbtune` table has the following definition.

Column Name	Data Type	Length	Allows Null
keyword	varchar	32	no
parameter_name	varchar	32	no
config_string	varchar	2048	yes

After creating the `SDE_dbtune` table, the ArcSDE setup populates the table with the contents of the `dbtune.sde` file, located in the `%SDEHOME%\etc` directory. If an `SDE_dbtune` table already exists, the ArcSDE setup may update some of its contents.

The keyword field stores keywords referenced when creating data in the database. Within each keyword, there are a number of parameters, each stored in the `parameter_name` field. Each parameter has a configuration string associated with it, stored in the `config_string` field.

You use the `SDE_dbtune` table for these purposes:

1. Load tables and indexes onto specific filegroups.
2. Specify an index created as either clustered or non-clustered.
3. Dictate the percentage to which an index page is filled.
4. Determine the amount of binary data to be inlined into data pages.
5. Control cross database query behavior – whether or not to filter feature classes and rasters at database boundaries.
6. Add a new keyword to apply different settings to data creation.

For most databases, the defaults listed in the SDE_dbtune table should be acceptable, except for the CROSS_DB_QUERY_FILTER parameter. This parameter prohibits ArcSDE from returning lists of all feature classes and rasters in all databases that the connecting user has permission to or owns.

How it Works

When you load data, create metadata server tables, or build geodatabase complex features, ArcSDE reads values from the SDE_dbtune table to determine how to create tables and indexes. There is a minor exception to this, CROSS_DB_QUERY_FILTER, which is read upon connection. The values read from SDE_dbtune are appended to CREATE TABLE and CREATE INDEX statements.

For example, you use the sderaster command to load a raster into an ArcSDE spatial database but have made some configuration changes that you want to apply to the raster when it is created. You have added a keyword, RASTERS, to the SDE_dbtune table that instructs ArcSDE to fill the block table's composite index to 95% capacity. You reference that keyword with the -k switch of the sderaster command. For example:

```
sderaster -o import -l ColoFloodPlain,image -g -f cfp0456.tif -k RASTERS ...
```

The RASTERS keyword has a parameter, BLK_INDEX_COMPOSITE, that allows you to specify how much to fill the index pages of the composite index on the raster's block table, and on which filegroup to create this index. The entry in the SDE_dbtune table would look like this:

Keyword	Parameter_name	Config_string
RASTERS	BLK_INDEX_COMPOSITE	WITH FILLFACTOR=95 ON RAS_FG

When you execute the sderaster command above, the loader reads the -k RASTERS keyword, and creates the block table's composite index with a fillfactor of 95% on a filegroup called "RAS_FG." T-SQL data definition language (DDL) would look like this:

```
CREATE UNIQUE CLUSTERED
    INDEX [SDE_blk_4_pk] ON [dbo].[SDE_blk_4] ([rasterband_id],
[rdd_factor], [row_nbr], [col_nbr])
WITH
    FILLFACTOR = 95
ON [RAS_FG]
```

Parameter_names not listed under the RASTERS keyword are read from the DEFAULTS keyword. Parameter_names will be redundant in this table; each keyword can specify a parameter already listed under DEFAULTS.

The rules for the SDE_dbtune table are:

1. SDE_dbtune is used to control data placement, index fill factors, index clustering, binary data inlining, and cross database querying ability.
2. If you don't specify a keyword when you load data, you use settings defined under DEFAULTS.
3. If you supply an incomplete keyword, missing parameters and config_strings are read from the DEFAULTS keyword.
4. You are not required to make any additional edits to SDE_dbtune to make the software work.
5. You can add new keywords with any valid name. You must use the parameter_name and config_string values supplied in this chapter. You cannot invent new parameter_names and config_strings.

6. You can modify existing config_string values.

Editing the SDE_dbtune table

Administrators edit the SDE_dbtune table to add keywords and change config_strings. You can edit the contents of the SDE_dbtune table with the SQL Server Enterprise Manager, sql statements, or the sdedbtune command.

To edit the SDE_dbtune table with the SQL Server Enterprise Manager, find the table in the tables list within the database node. Right click the table and select “open table,” then “return all rows.” You can also select the table, click the Action menu and select “open table” followed by “return all rows.” Make changes to config_strings by clicking within a cell. Press return to commit the change. To insert new dbtune records, scroll to the end of the table and click in the cell to the right of the asterisk. Keyword values cannot be null.

Using t-sql statements is more complex as you’ll have to be familiar with the insert, update and delete statements. These topics are beyond the scope of this book. Consult the Microsoft SQL Server Books Online (BOL) for more information.

The sdedbtune command exports the contents of the SDE_dbtune table to a file. The file can then be edited and re-imported into the SDE_dbtune table. In the following example, the SDE_dbtune table is exported, edited, and then re-imported into the source table.

```
C:\>sdedbtune -o export -f dbtune.out -s theodolite -D sde -u sde

C:\>write %sdehome%\etc\dbtune.sde

C:\>sdedbtune -o import -f dbtune.out -s theodolite -D sde -u sde
ArcSDE 8.1 Build 664 Tue Dec 26 22:32:22 PST 2000
Attribute Administration Utility
-----
Import SDE_dbtune Table. Are you sure? (Y/N): y

Successfully imported from file SDEHOME\etc\dbtune.out
```

The sdedbtune command always exports the file to the etc directory of the ArcSDE home directory.

Understanding the SDE_dbtune keywords

The SDE_dbtune table is grouped by keyword. At ArcSDE 8.3, the following keywords are included by default:

```
Select distinct(keyword) from sde.sde_dbtune

DATA_DICTIONARY
DEFAULTS
IMS_METADATA
IMS_METADATAARELATIONSHIPS
IMS_METADATAATAGS
IMS_METADATAATHUMBNAILS
IMS_METADATAUSERS
IMS_METADATAVALUES
IMS_METADATAWORDINDEX
IMS_METADATAWORDS
```

```
LOGFILE_DEFAULTS
NETWORK_DEFAULTS
NETWORK_DEFAULTS::DESC
NETWORK_DEFAULTS::NETWORK
SURVEY_MULTI_BINARY
TOPOLOGY_DEFAULTS
TOPOLOGY_DEFAULTS::DIRTYAREAS
```

You reference a keyword when loading data with an ArcSDE command tools, ArcToolbox, ArcCatalog, or create ArcIMS metadata server documents. When you build ArcGIS complex features like geometric networks, you can also reference a keyword.

DEFAULTS Keyword

The DEFAULTS keyword has 51 parameter_names and config_strings. Whenever you create an object within an ArcSDE database and don't specify an SDE_dbtune keyword, the DEFAULTS keyword supplies all table and index creation parameters. If you modify config_string values for the DEFAULTS keyword, those customizations apply to all new tables and indexes created when no dbtune keyword is specified. If you customize other keywords, but do not supply all parameter_names or config_strings, missing values will be read from the DEFAULTS keyword list.

When your application connects to a spatial database, ArcSDE reads the value of the CROSS_DB_QUERY_FILTER parameter_name in the DEFAULTS keyword to determine whether or not ArcSDE should return a list of all feature classes and rasters that the connecting user can access in any database or only the presently connected database. To enable CROSS_DB_QUERY_FILTER, set its config_string to 1.

SDE_dbtune sample showing some records of the DEFAULTS keyword

KEYWORD	PARAMETER_NAME	CONFIG_STRING
DEFAULTS	A_CLUSTER_USER	0
DEFAULTS	A_INDEX_ROWID	WITH FILLFACTOR = 75

DATA_DICTIONARY keyword

When the ArcSDE system tables are created during setup, the ArcSDE and geodatabase system tables and indexes are created with the storage parameters of the DATA_DICTIONARY keyword. To influence where and how the DATA_DICTIONARY tables are created, you have two options:

1. Drop any existing ArcSDE system tables (be careful), edit the dbtune.sde file in %SDEHOME%etc, and then run sdesetupmssql -o install.
2. Create a custom dbtune.sde file, run the ArcSDE Post-Installation utility, and browse to it during the "ArcSDE configuration files" portion of the setup.

To alter placement of existing tables, use the Enterprise Manager or T-SQL to add files and filegroups. See the SQL Server Books Online (BOL) for more information.

If your site makes extensive use of a multiversioned database, consider separating the SDE_MVTABLES_MODIFIED, SDE_STATE_LINEAGES, and SDE_STATES tables and indexes into their own filegroups to reduce I/O contention. If the dbtune.sde file does not contain the DATA_DICTIONARY keyword, or if any

of the required parameters are missing from the keyword, the following records will be inserted into the DATA_DICTIONARY when the table is created.

SDE_dbtune sample showing some records of the DATA_DICTIONARY keyword

KEYWORD	PARAMETER_NAME	CONFIG_STRING
DATA_DICTIONARY	B_INDEX_USER	WITH FILLFACTOR = 75
DATA_DICTIONARY	B_STORAGE	

LOGFILE keywords

SDE Logfiles are used by ArcSDE to maintain temporary and persistent sets of selected records. There are two logfile tables per user, SDE_Logfiles and SDE_Logfile_data. Both tables are created during the first connection a user makes, or recreated if they have been deleted the next time a user connects.

You can create keywords for a user's logfiles by appending the username to the LOGFILE_<username> keyword. The next time that user connects to an ArcSDE data store, logfiles will be created according to the parameter_names and config_strings specified in that keyword. Missing parameter_names will default to LOGFILE_DEFAULTS.

Most installations of ArcSDE will function well using the LOGFILE_DEFAULTS parameters. However, for applications that make heavy use of logfiles, such as ArcInfo, ArcEditor™, ArcIMS® MetaData Server, and ArcView, it may help scalability to spread the logfiles across separate filegroups. ArcGIS clients write logfiles whenever a selection set exceeds 100 records. Best practice, however, would be to change the client side Geodatabase registry setting for SelectionThreshold, boosting it to over 1000 records.

SDE_dbtune sample showing some records of the LOGFILE_DEFAULTS keyword

KEYWORD	PARAMETER_NAME	CONFIG_STRING
LOGFILE_DEFAULTS	LD_CLUSTER_ROWID	1
LOGFILE_DEFAULTS	LD_INDEX_ALL	WITH FILLFACTOR = 75

IMS_METADATA Keywords

The IMS METADATA keywords control storage of IMS Metadata Server tables. These keywords are a standard part of the dbtune table but are only meaningful at ArcSDE 8.2 and beyond. There are eight keywords with 64 total records.

ArcIMS Metadata Server Keywords

```
IMS_METADATA
IMS_METADATAARELATIONSHIPS
IMS_METADATAATAGS
IMS_METADATAATHUMBNAI
IMS_METADATAUSERS
IMS_METADATAVALUES
IMS_METADATAWORDINDEX
IMS_METADATAWORDS
```


The ArcIMS Metadata Server data model has one sde layer and nine tables. The layer, thumbnails, wordindex, words, and tags tables can grow to a large size. The data supplied by the Gazetteer service has 350000 records for the layer and thumbnails tables, 500000 for the wordindex table, and 55000 for the values table. Make sure that when you setup an ArcIMS metadata service or gazetteer service, your tables and indexes have enough room to grow.

If the keywords are not present in the dbtune file when it is imported into the SDE_dbtune table, ArcSDE applies software defaults. For more information about installing IMS Metadata and the associated tables and indexes, refer to ArcIMS Metadata Server documentation.

SDE_dbtune sample showing some records of the ims_metadata keywords

KEYWORD	PARAMETER_NAME	CONFIG_STRING
IMS_METADATA	S_INDEX_SP_FID	WITH FILLFACTOR = 75
IMS_METADATA	S_STORAGE	

The NETWORK_DEFAULTS keyword

The NETWORK_DEFAULTS keyword contains the default parameters for the ArcGIS network class. If no keyword is supplied upon creation of a geometric network, parameters are read from this keyword. If a keyword containing incomplete parameter_names is used during network creation, missing parameter names will be supplied from the DEFAULTS keyword.

A geometric network will create several new tables within the database. The largest table will generally be the N_<id>_DESC table so make sure its corresponding SDE_dbtune keyword, NETWORK_DEFAULTS::DESC either explicitly lists or defaults to a filegroup that will accommodate large table and index growth.

Network class composite keywords

The composite keyword is a unique type of keyword designed to accommodate the tables of the ArcGIS network class. The network table's size variation requires a keyword that provides configuration parameters for both large and small tables. Typically, the network descriptions table is very large in comparison with the others. To accommodate the vast difference in size of the network tables, the network composite keyword is subdivided into elements. A network composite keyword has three elements: the parent element defines the general characteristic of the keyword and the junctions feature class, the description element defines the configuration of the DESCRIPTIONS table and its indexes, and the network element defines the configuration of the remaining network tables and their indexes. The parent element does not have a suffix, and its keyword looks like any other keyword. The description element is demarcated by the addition of the ::DESC suffix to the parent element's keyword, and the network element is demarcated by the addition of the ::NETWORK suffix to the parent element's keyword.

SDE_dbtune sample showing some records from the NETWORK_DEFAULTS keyword

KEYWORD	PARAMETER_NAME	CONFIG_STRING
NETWORK_DEFAULTS::DESC	D_INDEX_DELETED_AT	WITH FILLFACTOR = 75
NETWORK_DEFAULTS::DESC	D_STORAGE	

The TOPOLOGY keyword

The TOPOLOGY keyword controls the storage of topology tables, which are named POINTERRORS, LINEERRORS, POLYERRORS and DIRTYAREAS. An SDE instance must have a valid topology keyword in the dbtune table, or topology will not be built.

The DIRTYAREAS table maintains information on areas within a layer that have been changed. Because it tracks versions, data will be inserted or updated but not deleted during normal use. The DIRTYAREAS table will reduce in size only when database versions get compressed.

Because the DIRTYAREAS table is much more active than the remaining topology tables, the TOPOLOGY keyword may be compound. You may specify the DIRTYAREAS suffix to list configuration string to be used to create the topology tables.

SDE_dbtune sample showing some records of the TOPOLOGY_DEFAULTS keyword

KEYWORD	PARAMETER_NAME	CONFIG_STRING
TOPOLOGY_DEFAULTS	B_INDEX_USER	WITH FILLFACTOR = 75
TOPOLOGY_DEFAULTS	B_STORAGE	

Parameter_names and Config_Strings

The previous section listed the default keywords but did not explain parameter_names and their config_strings. Parameter_names and config_strings are used by ArcSDE to identify:

1. Where a table or index is created (filegroup).
2. Whether or not to cluster an index.
3. How much to fill each index page (FILLFACTOR).
4. How much binary data should be stored inline to a data page (TEXT_IN_ROW).
5. Whether to allow for cross database query support (CROSS_DB_QUERY_FILTER).

Parameter_names identify the object to be created, indexed, fillfactored, or inlined, while config_strings identify how that object will be created, indexed, fillfactored, or inlined. Parameter_names are fixed; you cannot invent new parameter_names. Likewise, Config_Strings accept only certain sql strings. These strings are appended to sql CREATE TABLE and CREATE INDEX statements.

Parameters and config_strings fall into these categories:

Parameter_name Category	Config_String Values Types	Parameter_name Example	Config_String Example
Cluster an index	1 = yes; 0 = no	A_CLUSTER_SHAPE	1
Index storage and fillfactor	WITH FILLFACTOR=75 ON <filegroup>	A_INDEX_SHAPE	WITH FILLFACTOR=75 ON mv_fg
Table Storage	ON <filegroup>	A_STORAGE	ON mv_fg
TEXT_IN_ROW	Inline up to 256 bytes (default) in the data page	F_TEXT_IN_ROW	256
list feature class & raster names in all or presently connected db	1 = yes; 0 = no	CROSS_DB_QUERY_FILTER	0
Expose a keyword in a gui (UI_TEXT)	Descriptive text to appear in gui	UI_TEXT	RASTERS
Comment	Descriptive text	COMMENT	Raster Info

Table and Index Parameters and Config_strings

The SDE_dbtune table is grouped by keyword. Therefore, parameter_names are non-unique and are re-used many times throughout the table.

Adds Table (Multi-version table) Parameters

An “Adds table” is a multi-versioned table that stores edits made against a feature class in a multi-versioned database. It is almost identical in structure to the business table but has additional columns to track state ids.

Parameter_name	Default Config_string	Note	Example Config_String
A_CLUSTER_ROWID	0	Create non-clustered A table rowid index	0
A_CLUSTER_SHAPE	1	Create clustered A table spatial col (shape) index	1
A_CLUSTER_STATEID	0	Create non-clustered A table stateid index	0
A_CLUSTER_USER	0	Create non-clustered A table user defined index (if exists)	0
A_INDEX_ROWID	WITH FILLFACTOR = 75	Fill the index pages for rowid index to 75%	with fillfactor = 75 on ras_fg
A_INDEX_SHAPE	WITH FILLFACTOR = 75	Fill the shape index pages to 75%	with fillfactor = 75 on ras_fg
A_INDEX_STATEID	WITH FILLFACTOR = 75	Fill the stateid index pages to 75%	with fillfactor = 75 on ras_fg
A_INDEX_USER	WITH FILLFACTOR = 75	Fill any user defined indexes to 75%	with fillfactor = 75 on ras_fg
A_STORAGE		Specify a filegroup with on <fg>	on ras_fg

Business Table Parameters

The business table is the attribute table of a layer or feature class. It contains attribute information in addition to a spatial column (of type int) used to join to the feature table.

Parameter_name	Default Config_String	Note	Example Config_String
B_CLUSTER_ROWID	0	Create non-clustered rowid index	0
B_CLUSTER_SHAPE	1	Create clustered spatial col (shape) index	1
B_CLUSTER_USER	0	Create non-clustered user defined index	0
B_INDEX_ROWID	WITH FILLFACTOR = 75	Fill rowid index pages to 75%;	with fillfactor = 75 on bus_fg
B_INDEX_SHAPE	WITH FILLFACTOR = 75	Fill shape index pages to 75%;	with fillfactor = 75 on bus_fg
B_INDEX_USER	WITH FILLFACTOR = 75	Fill a user defined index page to 75%	with fillfactor = 75 on bus_fg
B_TEXT_IN_ROW	256	store up to 256 bytes directly int the data page for binary type data	256
B_STORAGE		Specify a filegroup with ON <filegroup name>	on bus_fg

Deletes Table Parameters

The deletes table is used to track modifications against multi-versioned tables.

Parameter_name	Default Config_string	Note	Example Config_String
D_CLUSTER_ALL	0	Create non-clustered composite index	0
D_CLUSTER_DELETED_AT	1	Create clustered deleted_at column index	1
D_INDEX_ALL	WITH FILLFACTOR = 75	Fill composite index pages to 75%	WITH FILLFACTOR = 75 on d_fg
D_INDEX_DELETED_AT	WITH FILLFACTOR = 75	Fill deleted_at index pages to 75%	WITH FILLFACTOR = 75 on d_fg
D_STORAGE		Specify a filegroup with ON <filegroup name>	on d_fg

Feature Table Parameters

The feature table stores each shape's extent and geometry. It will also contain records from multi-versioned inserts and updates.

Parameter_name	Default Config_string	Note	Example Config_String
F_CLUSTER_FID	1	Cluster FID column index	1
F_INDEX_AREA	WITH FILLFACTOR = 75	Fill area index pages to 75%	WITH FILLFACTOR=75 ON f_fg
F_INDEX_FID	WITH FILLFACTOR = 75	Fill the FID index pages to 75%	WITH FILLFACTOR=75 ON f_fg
F_INDEX_LEN	WITH FILLFACTOR = 75	Fill the len index pages to 75%	WITH FILLFACTOR=75 ON f_fg
F_STORAGE		Specify a filegroup with on <filegroup>	on f_fg
F_TEXT_IN_ROW	256	Store up to 256 bytes of binary data in the data page	256

Spatial Index Table Parameters

The spatial index is a grid that overlays features and is used to identify features to fetch. The bounding box of spatial query is overlaid against the spatial index table to select candidate shapes satisfying the query.

Parameter_name	Config_string	Note	Example Config_String
S_CLUSTER_ALL	1	Create clustered composite index	1
S_CLUSTER_SP_FID	0	Create non-clustered sp_fid index	0
S_INDEX_ALL	WITH FILLFACTOR = 75	Fill composite index pages to 75%	WITH FILLFACTOR=75 ON sp_fg
S_INDEX_SP_FID	WITH FILLFACTOR = 75	Fill sp_fid index pages to 75%	WITH FILLFACTOR=75 ON sp_fg
S_STORAGE		Specify a filegroup with on <filegroup>	on sp_fg

Data Dictionary Keyword Parameter_names

These parameters govern how the ArcSDE some metadata tables will be created. The data dictionary keyword also contains Business table parameters. These parameters are listed above in “Business Table Parameters.”

Parameter_name	Config_string	Note	Example Config_String
MVTABLES_MODIFIED_INDEX	WITH FILLFACTOR = 75	Fill mvtables_modified index to 75%	WITH FILLFACTOR = 75 on dd_fg
MVTABLES_MODIFIED_TABLE		Specify a filegroup for mv_tables_modified	on dd_fg
STATE_LINEAGES_INDEX	WITH FILLFACTOR = 75	Fill state_lineages index pages to 75%	WITH FILLFACTOR = 75 on dd_fg
STATE_LINEAGES_TABLE		Specify a filegroup for state_lineages table	on dd_fg
STATES_INDEX	WITH FILLFACTOR = 75	Fill states table index pages to 75%	WITH FILLFACTOR = 75 on dd_fg
STATES_TABLE		Specify filegroup for states table	on dd_fg
VERSIONS_INDEX	WITH FILLFACTOR = 75	fill versions table index pages to 75%	WITH FILLFACTOR = 75 on dd_fg
VERSIONS_TABLE		specify a filegroup for versions table	on dd_fg

Logfile Tables Parameters

As described previously, the sde_logfile and sde_logfile_data tables hold selected sets of records.

Parameter_name	Default	Config_string	Note	Example Config String
LD_CLUSTER_ALL	0		Create non-clustered composite index on Logfile_Data table	0
LD_CLUSTER_DATA_ID	0		Create non-clustered data_id index on Logfile_data table	0
LD_CLUSTER_ROWID	1		Create clustered rowid index on Logfile_data table	1
LD_INDEX_ALL	WITH FILLFACTOR = 75		Fill composite index pages to 75%	WITH FILLFACTOR = 75 on log_fg
LD_INDEX_DATA_ID	WITH FILLFACTOR = 75		Fill data_id index pages to 75%	WITH FILLFACTOR = 75 on log_fg
LD_INDEX_ROWID	WITH FILLFACTOR = 75		Fill rowid index pages to 75%	WITH FILLFACTOR = 75 on log_fg
LD_STORAGE			Specify filegroup Logfile_data table	on log_fg
LF_CLUSTER_DATA_ID	0		Create non-clustered data_id index on logfiles table	0
LF_CLUSTER_ID	0		Create non-clustered id index on the logfiles table clustered	0
LF_CLUSTER_NAME	0		Create non-clustered name index on logfiles table	0
LF_INDEX_DATA_ID	WITH FILLFACTOR = 75		Fill data_id index pages to 75%	WITH FILLFACTOR = 75 on log_fg
LF_INDEX_ID	WITH FILLFACTOR = 75		Fill id index pages to 75%	WITH FILLFACTOR = 75 on log_fg
LF_INDEX_NAME	WITH FILLFACTOR = 75		Fill name index pages to 75%	WITH FILLFACTOR = 75 on log_fg
LF_STORAGE			Specify filegroup for Logfiles table	on log_fg

IMS_METADATA, NETWORK_DEFAULTS, and TOPOLOGY Parameters

The parameter_names for these keywords work the same way as the parameters and config_strings of regular layers or business tables. You can customize the index fillfactors of business, feature, and spatial index tables and create the data on filegroups. You should not change the clustering parameters.

Raster table parameters

Rasters in ArcSDE are stored as 5 separate tables. Of these tables, only one will get big, the BLK table. Set the BND_CLUSTER_COMPOSITE config_string to 1.

Parameter_name	Config_string	Note	Example Config_string
AUX_CLUSTER_COMPOSITE	1	Cluster composite index on the aux table	
AUX_INDEX_COMPOSITE	WITH FILLFACTOR = 75	Fill the AUX table index pages to 75%	WITH FILLFACTOR = 75 on ras_fg
AUX_STORAGE		Specify a filegroup for aux table	on ras_fg
BLK_CLUSTER_COMPOSITE	1	Cluster BLK table composite index	1
BLK_INDEX_COMPOSITE	WITH FILLFACTOR = 75	Fill the BLK table composite index pages	WITH FILLFACTOR = 75 on ras_fg
BLK_TEXT_IN_ROW	256	Store up to 256 bytes in data page of binary data	
BLK_STORAGE		Specify a filegroup for blk table	on ras_fg
BND_CLUSTER_COMPOSITE	0	Create non-clustered BND table composite	0
BND_CLUSTER_ID	0	Create non-cluster BND table id index	0
BND_INDEX_COMPOSITE	WITH FILLFACTOR = 75	Fill the BLK table's composite index pages to 75%	WITH FILLFACTOR = 75 on ras_fg
BND_INDEX_ID	WITH FILLFACTOR = 75	Fill the BLK table's id index pages to 75%	WITH FILLFACTOR = 75 on ras_fg
BND_STORAGE		Specify a filegroup for BND table	on ras_fg
RAS_CLUSTER_ID	1	Cluster RAS table id	1
RAS_INDEX_ID	WITH FILLFACTOR = 75	Fill the RAS table's id index pages to 75%	WITH FILLFACTOR = 75 on ras_fg
RAS_STORAGE		Specify a filegroup for RAS table	on ras_fg

UI_TEXT Parameters

Any parameters listed as “UI_TEXT” represent dbtune keywords visible to GUIs. Tools like the Network Builder in ArcCatalog expose dbtune keywords for data loading and modeling operations.

Text in row

SQL Server 2000 supports the ability to store image, text and ntext values directly in a data row. Previous versions of SQL Server would store the images types in image data pages. The advantage of “inlining” or storing some image types in a data row is that you may speed fetching of your data since queries will traverse fewer pages to find your data. ArcSDE uses an image type column to store its geometry.

The SQL Server and ArcSDE default of for text in row is 256 bytes, but can span from 72 to 7000 bytes. Be forewarned that the more image data you store in a row, the larger your table will become, thus eventually slowing your queries as more pages must be traversed. Larger raster data should not be stored in row. Points do not need to be stored in row as spatial queries against them reconstruct the features from their envelopes instead of reading from the image type.

To enable a table or feature class to store its geometry, rasters, regular images, text or ntext “in row” reference a keyword containing the appropriate text in row values:

Keyword	Parameter_name	Config_string
DEFAULTS	B_TEXT_IN_ROW	256
DEFAULTS	F_TEXT_IN_ROW	256
IMS_METADATA	B_TEXT_IN_ROW	256
IMS_METADATA	F_TEXT_IN_ROW	256
IMS_METADATAHUMBNAILS	B_TEXT_IN_ROW	256
NETWORK_DEFAULTS	B_TEXT_IN_ROW	256
NETWORK_DEFAULTS	F_TEXT_IN_ROW	256
NETWORK_DEFAULTS::DESC	B_TEXT_IN_ROW	256
NETWORK_DEFAULTS::NETWORK	B_TEXT_IN_ROW	256

Picking a text_in_row size

Try to inline up to 80% of your vector data. You can ascertain this size by running this query:

```
select top 20 percent (datalength(points)) from boris.fl2 order by
datalength(points) desc
```

Pick the smallest value from this result set. Raster data is usually much larger. If you attempt to inline too much raster or vector data, you'll end up creating too many extents and data pages. Generally, the default value(256) should suffice.

Fill factor parameters

The "WITH FILLFACTOR = 75" parameter specifies how full each page in the leaf level of an index should be. SQL Server uses a default value of 0, which means that the leaf pages of an index are almost full but that the nonleaf pages have room for at least two more rows. If the fillfactor is 100, all pages are completely full. With a fill factor of 75, each clustered index page starts 75 percent full. Subsequent inserts and updates to that data add to the index page. When the page hits 100%, it is full. Any subsequent insert or update to data in that page will split the page.

Use FILLFACTOR to balance full index pages and page splits. When a page is split, SQL Server moves approximately 50% of the data in the split page to a new page, most likely allocated from a different extent. Page splits will fragment your tables and compromise performance. Setting Fillfactor too low creates too many data pages and extents to traverse in a query, thus negatively impacting performance.

Decision Criteria for choosing a FILLFACTOR

1. Is your data read-only? Will it never be edited? If so, set all FILLFACTORS on your data to 100.
2. Will your data be updated frequently? If so, stay with the defaults.
3. Will your data incur sporadic updates? Pick a range between 75-95% based on how often you want to defragment your tables.

Clustered Indexes and their Tables

Clustered indexes store tabular data at their leaf nodes. The data pages at the clustered index leaf level derive their order from the clustered index key value. This has one important consequence with regard to the SDE_dbtune table: **You cannot separate a table from its clustered index.**

For example, you specify that a feature table's fid index be created on the FeatIdx filegroup while the feature table should be stored on the Feat filegroup. The fid index is created as clustered. SDE_dbtune might look like this:

Keyword	Parameter_name	Config_string
DEFAULTS	F_INDEX_FID	WITH FILLFACTOR=90 ON FEATIDX
DEFAULTS	F_STORAGE	ON FEAT

In this case, both the feature table and feature table's index will reside on the FeatIdx filegroup. The feature table is created first, then a primary key constraint is applied to the FID column. The constraint creates a clustered index on the fid column and references the FEATIDX filegroup in this statement:

```
Alter Table features.dbo.f4 add constraint f4_pk primary key clustered
(fid) with fillfactor=75 on FEATIDX
```

Another case would be if you were to specify a filegroup for the feature table but not for the feature table's index.

Keyword	Parameter_name	Config_string
DEFAULTS	F_INDEX_FID	WITH FILLFACTOR=90
DEFAULTS	F_STORAGE	ON FEAT

In this case, both the feature table and clustered index on the fid column will reside on the Feat filegroup. This occurs because the table is created first and when the alter table statement is applied, no ON statement is appended because no such string is listed in the config_string column above.

```
Alter table features.dbo.f5 add constraint f5_pk primary key clustered
(fid) with FILLFACTOR=75
```

CHAPTER 4

Managing tables, feature classes, rasters, and views

A fundamental part of any database is creating and loading the tables. Tables with spatial columns are called standalone feature classes. Attribute only (nonspatial) tables are also an important part of any database. This chapter will describe the table and feature creation and loading process and provide some examples demonstrating how to populate your ArcSDE database.

Setting up the dbtune table

Previous versions of ArcSDE and Spatial Database Engine™ (SDE®) used the dbtune.sde file to identify data spaces for your data. Loading your data onto specified data spaces allows you to distribute I/O in multiuser configurations. The basic idea is to spread your data and indexes across many disks so that disk spindles can find it. Placing all data in one location on a single disk will create a “hotspot.” A hotspot is a location on a disk of frequently accessed data by multiple readers or writers.

ArcSDE 8.3 pushes the dbtune file into a database table. This allows for all types of connections, three-tier or two-tier, to access the information. Dbtune tables are composed of keywords, parameters, and configuration parameters. Keywords identify specific loading parameters. Parameters represent ArcSDE objects such as tables and indexes. Configuration parameters specify how something should be loaded.

For further information on the dbtune table, see Chapter 3, ‘Configuring DBTUNE storage parameters’.

Data creation

There are numerous applications that can create and load data within an ArcSDE SQL Server database. These include:

1. ArcSDE administration commands located in the bin directory of SDEHOME:
 - € sdelayer manages feature classes including creating a schema
 - € sdetable manages tables including creating a schema
 - € sdeimport takes an existing sdeexport file and loads the data into the database
 - € shp2sde loads an ESRI® shapefile into the database
 - € cov2sde loads a coverage, Map LIBRARIAN™ layer, or ArcStorm™ layer into the database

- € tbl2sde loads an attribute-only dBASE or INFO file into the database
- € sdeggroup—a specialty feature class creation command that provides a way to “group” or tile features from an existing feature class into another feature class that displays background or reference data. The generated feature classes are used for rapid display of a large amount of geometry data. The attribute information is not retained, and spatial searches cannot be performed on these feature classes.

These are all run from the operating system prompt. Command references for these tools are in the ArcSDE developer help.

Other applications include:

ArcGIS Desktop—use ArcCatalog or ArcToolbox™ to manage and populate your database.

ArcInfo Workstation—use the Defined Layer interface to create and populate the database.

ArcView GIS 3.2—use the Database Access extension.

MapObjects®—custom COM applications can be built to create and populate databases.

ArcSDE CAD Client extension—for AutoCAD and MicroStation users.

Other third party applications built with either C or Java APIs.

This document focuses primarily on the ArcSDE administration tools but does provide some ArcGIS Desktop examples as well. In general, most people prefer an easy-to-use graphic user interface like the one found in ArcGIS 8.3 Desktop. For details on how to use ArcCatalog or ArcToolbox (another desktop data loading tool), please refer to the ArcGIS 8.3 books:

Know your data

Before you begin loading your data, get to know it. SQL Server 2000 has approximately 174 reserved words and 135 future keywords, words that may be listed as reserved in future releases. The use of reserved words for table and column names is not supported. See the SQL Server Books Online for more information.

The use of nonstandard identifiers is also unsupported. Such an identifier could be naming a table “3d.” You cannot begin an ArcSDE tablename with a number. Nor can you use pound signs or symbols. The maximum length of an ArcSDE table is 32 characters.

Watch the length and naming of your columns as well. Do not employ column names with spaces in them. Long column names can also cause problems if they exceed the SQL Server accepted range. Finally, be aware of your data ranges, especially with date data types. SQL Server cannot accept datetime values preceding January 1, 1753. If you attempt to load a table into SQL Server through ArcSDE with values preceding this range, the load will fail.

Creating and populating a feature class

The general process involved with creating and loading a feature class is:

1. Create the business table.
2. Add the feature class to the ArcSDE system tables.

3. Switch the feature class to load_only_io mode (optional step to improve bulk data loading performance. It is OK to leave the feature class in normal_io mode to load data.).
4. Insert the records (load data).
5. Switch the feature class to normal_io mode (builds the indexes).
6. Version the data (optional).
7. Grant permissions on the data (optional).

In the following sections, this process is discussed in more detail and illustrated with some examples of ArcSDE administration commands usage and ArcInfo data-loading utilities through the ArcCatalog and ArcToolbox interfaces.

Creating a feature class “from scratch”

There are two basic ways to create a feature class. You can create a feature class from scratch (generally regarded as more effort), or you can create a feature class from existing data such as a coverage or ESRI shapefile. Both methods are reviewed below with the “from scratch” method being first.

Creating a business table

You may create a business table with either the SQL CREATE TABLE statement or with the ArcSDE `shtable` command. The `shtable` command allows you to include a `dbtune` configuration keyword which contains the storage parameters of the table.

Although the table may be up to 256 columns, ArcSDE requires that only one of those columns be defined as a spatial column.

In this example, the `shtable` command is used to create the ‘roads’ business table.

```
shtable -o create -t roads -d 'road_id integer, name string(32), shape
integer' -k roads -u beetle -p bug
```

The table is created using the `dbtune` configuration keyword (-k) ‘roads’ by user beetle.

The same table could be created with a SQL CREATE TABLE statement using the SQL Server Query Analyzer.

```
create table roads
(road_id      integer,
 name        varchar(32) ,
 shape       integer)
on beetle_data
```

At this point you have created a table in the database. ArcSDE does not yet recognize it as a feature class. The next step is to add a feature class.

Adding a feature class

After creating a business table, you must add an entry for it in the ArcSDE system tables before the ArcSDE server can reference it. Use the `sdelay` command with the “-o add” operation to add the new feature class. This will add a spatial column, update the ArcSDE `SDE_layers` and `SDE_geometry_columns` tables, and add the required system table entries.

In the following example, the roads feature class is added to the ArcSDE database. Note that to add the feature class, the roads table name and the shape spatial column are combined to form a unique feature class reference. To understand the -e, -g, and -x options, refer to the `sdelay` command reference in the ArcSDE developer help.

```
sdelay -o add -l roads,shape -e 1+ -g 256,0,0 -x 0,0,100 -u beetle -p
bug -k roads
```

If the spatial column of the feature class is stored in either LONG RAW or BLOB compressed binary format, the feature and spatial index tables are created at this time. The feature class tables and indexes are stored according to the storage parameters of the **roads** configuration keywords in the DBTUNE table. Upon successful completion of the previous `sdetable` command—to create a table—and the `sdelay` command—to add the feature class to the ArcSDE system tables—you have an empty feature class in `normal_io` mode.

Switching to load-only mode

Switching the feature class to load-only mode drops the spatial index and makes the feature class unavailable to ArcSDE clients. Bulk loading data into the feature class in this state is much faster due to the elimination of index maintenance. Use the `sdelay` command to switch the feature class to load-only mode by specifying the “-o load_only_io” operation.

```
sdelay -o load_only_io -l roads,shape -u beetle -p bug
```

Note: A feature class, registered as multiversioned, cannot be placed in the load-only I/O mode. However, the grid size can be altered with the -o alter operation. The alter operation will apply an exclusive lock on the feature class, preventing all modifications by ArcInfo until the operation is complete.

Inserting records into the feature class

Once the empty feature class exists, the next step is to populate it with data. There are several ways to insert data into a feature class, but probably the easiest method is to convert an existing shapefile or coverage or import a previously exported ArcSDE `sdeexport` file directly into the feature class. A more “from scratch” method would be to add the data with an editor such as ArcMap.

In this first example, `shp2sde` is used with the `init` operation. The `init` operation is used on newly created feature classes or can be used on feature classes when you want to “overwrite” data that’s already there. Be careful when using the `init` operation on feature classes that already have data loaded in them. Here, the shapefile, “`rdshp`”, will be loaded into the feature class, “`roads`”. Note that the name of the spatial column (`shape` in this case) is included in the feature class (`-l`) argument.

```
shp2sde -o init -l roads,shape -f rdshp -u beetle -p bug
```

Similarly, we can also use the `cov2sde` command:

```
cov2sde -o init -l roads,shape -f rdcov -u beetle -p bug
```

Switching the table to normal_io mode

After data has been loaded into the feature class, you must switch the feature class to `normal_io` mode to re-create all indexes and make the feature class available to clients. For example:

```
sdelay -o normal_io -l roads,shape -u beetle -p bug
```

Now your feature class is ready for use by ArcSDE client applications.

Versioning your data

Optionally, you may enable your feature class as multiversioned. Versioning is a process that allows for multiple representations of your data without requiring duplication or copies of the data. ArcMap requires data to be multiversion enabled in order to edit it. For further information on versioning data, refer to the *Building a Geodatabase* book.

In this example, the feature class called ‘`states`’ will be registered as multiversioned using the `sdetable alter_reg` operation.

```
sdelayer -o alter_reg -t states -c ver_id -C SDE -V multi -k GEOMETRY_TYPE
```

Granting permissions on the data

Once you have the data loaded, it is often necessary for other users to have access to the data for update, query, insert, or delete operations. Initially, only the user who has created the business table has access to it. In order to make the data available to others, the owner of the data must grant permissions to other users. The owner can use the sdelayer command to grant permissions. Privileges can be granted to either another user or to a role.

In this example, a user called “beetle” gives a user called “spider” SELECT privileges on a feature class called ‘states’.

```
sdelayer -o grant -l states,feature -U spider -A SELECT -u beetle -p bug
```

The full list of -A keywords are:

SELECT. The user may query the selected object(s) data.

DELETE. The user may delete the selected object(s) data.

UPDATE. The user may modify the selected object(s) data.

INSERT. The user may add new data to the selected object(s) data.

If you include the -I grant option, you also grant the recipient the privilege of granting other users and roles the same privilege.

Creating and loading feature classes from existing data (recommended method)

This next section reviews how to create feature classes from existing data. This method is simpler since the creation and load process is completed at once.

Creating new feature classes and loading data all in one step

Each of the ArcSDE administration commands, shp2sde, cov2sde, and sdeimport, includes an “-o create” operation, which allows you to create a new feature class within the ArcSDE database. The create operation does all of the following:

- ≠ Creates the business table using the input data as the template for the schema
- ≠ Adds the feature class to the ArcSDE system tables
- ≠ Puts the feature class into load-only mode
- ≠ Inserts data into the feature class
- ≠ When all the records are inserted, puts the feature class into normal_io mode

shp2sde

The shp2sde command converts shapefiles into ArcSDE feature classes. The spatial column definition is read directly from the shapefile. You can use the shpinfo command to display the shapefile column definitions. In this example, the -k roadsys switch identifies a dbtune parameter.

```
shp2sde -o create -f rdshp -l roads,shape -k roadsys -u beetle -p bug
```

cov2sde

The cov2sde command converts ArcInfo coverages, ArcInfo Librarian™ library feature classes and ArcStorm library feature classes into ArcSDE feature classes. The create operation derives the spatial column definition from the coverage’s feature attribute table.

Use the ArcInfo describe command to display the ArcInfo data source column definitions.

In this example, an ArcStorm library, “roadlib”, is converted into the feature class, “roads”.

```
cov2sde -o create -l roads,shape -f roadlib,arcstorm -g 256,0,0 -x 0,0,100
-e 1+ -u beetle -p bug
```

sdeimport

The sdeimport command converts ArcSDE export files into ArcSDE feature classes. In this example, the roadexp ArcSDE export file is converted into the feature class ‘roads’.

```
sdeimport -o create -l roads,shape -f roadexp -u beetle -p bug
```

After using these commands to create and load data, you may optionally need to enable multiversioning on the feature class and modify the permissions of the feature class.

Appending data to an existing feature class

A common requirement for data management is to be able to append data to existing feature classes. All the data loading commands described thus far have a `-o append` operation for appending data. A feature class must exist prior to using the append operation. If the feature class is multiversioned, it must be in an “open” state. It is also advisable to change the feature class to load-only I/O mode and pause the spatial indexing operations before loading the data to improve the data loading performance. The spatial indexes will be re-created when the feature class is put back into normal I/O mode. Because the feature class has already been defined, the metadata already exists and is not altered by the append operation.

In the shp2sde example below, a previously created ‘roads’ feature class gets appended features from a shapefile, “rdshp2.” All existing features, loaded from the rdshp shapefile, will remain intact, and the feature class will simply be updated with the new features from the rdshp2 shapefile.

```
sdelayer -o load_only_io -l roads,shape -u beetle -p bug
shp2sde -o append -f rdshp2 -l roads,shape -u beetle -p bug
sdelayer -o normal_io -l roads,shape -u beetle -p bug
sdetable -o update_dbms_stats -t roads -u beetle -p bug
```

Note the last command in the sequence. The `update_dbms_stats` operation on `sdetable` updates the table and index statistics.

Creating and populating raster tables

Raster columns are created from the ArcGIS Desktop using ArcCatalog or ArcMap. To create a raster column, you will first need to convert the image file into a format acceptable to ArcSDE. Then after the image has been converted to the ESRI raster file format, you can convert it into a raster column.

For more information on creating raster columns using either ArcCatalog or ArcToolbox, refer to *Building a Geodatabase*. ArcCatalog and ArcToolbox have a raster to Geodatabase loader for loading rasters into ArcSDE.

Consult Chapter 3 ‘Configuring DBTUNE storage parameters,’ for dbtune specific parameters for loading rasters. ArcSDE 8.3 for SQL Server 2000 supports a `text_in_row` option for raster columns: `blk_text_in_row`.

Creating a Raster Image Catalog

An Image Catalog allows you to group many images by simply listing them in a table. ArcGIS clients such as ArcCatalog or ArcMap display the images as a group by reading the entries in the table. The table must contain five columns: Image, xmin, ymin, xmax, and ymax. The Image column contains a fully qualified image name, and the remaining four describe the extent of each image. The table does not have to be registered with either ArcSDE or the Geodatabase and does not have to be multiversioned. Here's an example:

Image	xmin	ymin	xmax	ymax
sde.DBO.A106	322169.2	4094888	411902.3	4151957
sde.DBO.E106	323452.5	4150579	412586.2	4207658
sde.DBO.A107	233002.1	4096472	323537.9	4154628
sde.DBO.E107	234912.9	4151847	324666.3	4210094
sde.DBO.A108	676712.7	4090675	770814.6	4154551
sde.DBO.E108	675530.7	4146368	769195.7	4210173
sde.DBO.A106_8	324599.5	4206098	413235.3	4263136
sde.DBO.E106_8	325340.9	4261470	413692.2	4318609
sde.DBO.A107_8	236598.8	4207336	325702	4265705
sde.DBO.E107_8	239302.9	4262352	327095.4	4321168
sde.DBO.A105_9	413273.8	4316886	500802.2	4372878

If you have less than nine images, then all the images will display. If you have more, your images will not display until you begin to zoom in. At ArcSDE 8.3, there is another type of image catalog, the embedded raster catalog. This type of catalog allows for multiple raster layers to be stored within the same series of tables. This option provides better performance than a traditional image catalog. More information about this is available in the white paper "Raster Data in ArcSDE 8.2."

Exporting data

As with importing data, there are also client applications that export data from ArcSDE as well. With ArcSDE, the following command line tools exist:

`sdeexport`—creates an ArcSDE export file to easily move feature class data between ArcSDE instances.

`sde2shp`—creates an ESRI shapefile from an ArcSDE feature class

`sde2cov`—creates a coverage from an ArcSDE feature class

`sde2tbl`—creates a dBASE or INFO file from an RDBMS table

Schema modification

There will be occasions where it is necessary to modify the schema of some tables. You may need to add or remove columns from a table. ArcCatalog offers an easy-to-use tool for this and other schema operations such as modifying the spatial index (grids) and adding and dropping column indexes. To drop a column with ArcSDE command tools, create an sde view of the layer without the column, export it, then reimport it. To add a column using the sde admin tools, create a table from your business table using a select into statement, then add your new column.

Using ArcInfo software's ArcCatalog and ArcToolbox applications

So far the discussion has focused on ArcSDE command line tools that create feature class schemas and load data into them. While robust, these commands can be daunting for the first-time user. In addition, if you are using ArcInfo Desktop, you may have to use ArcCatalog to create feature datasets and feature classes within those feature datasets to use specific ArcInfo 8 Desktop functionality. For that reason, we provide a glimpse of how to use ArcToolbox and ArcCatalog to load data. Please refer to the ArcInfo Digital Books on ArcCatalog, ArcToolbox, and the Geodatabase for a full discussion of these tools.

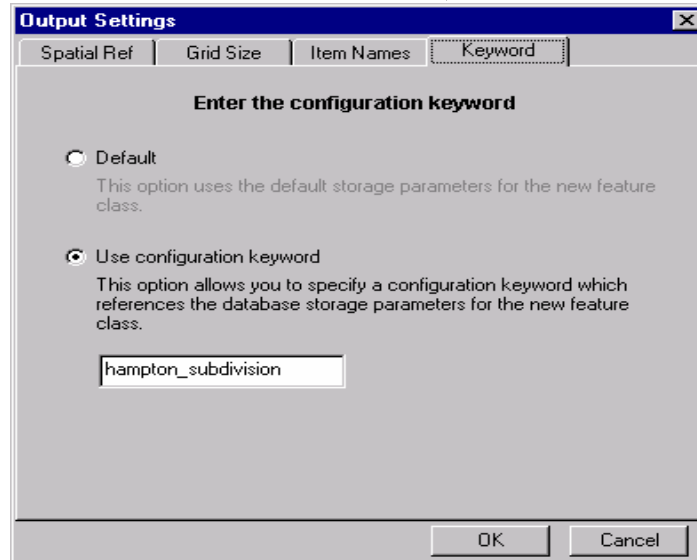
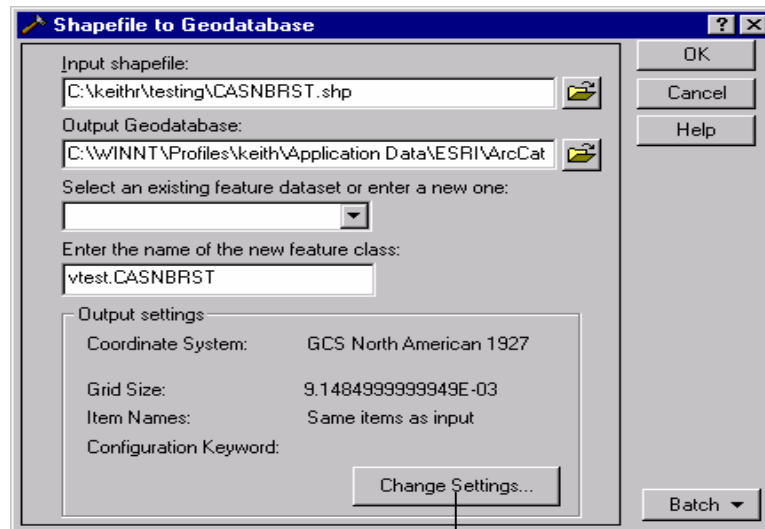
Loading data

You can convert ESRI shapefiles, coverages, Map LIBRARIAN layers, and ArcStorm layers into Geodatabase feature classes with the ArcInfo ArcToolbox and ArcCatalog applications. ArcToolbox provides a number of tools that enable you to convert data from one format to another.

ArcToolbox operations, such as the ArcSDE administration commands `shp2sde`, `cov2sde`, and `sdeimport`, accept configuration keywords. By using a configuration keyword from the `dbtune` table, you can employ customized data loading options such as fill factors of indexes, `text_in_row` sizes, or filegroups for storing tables or indexes.

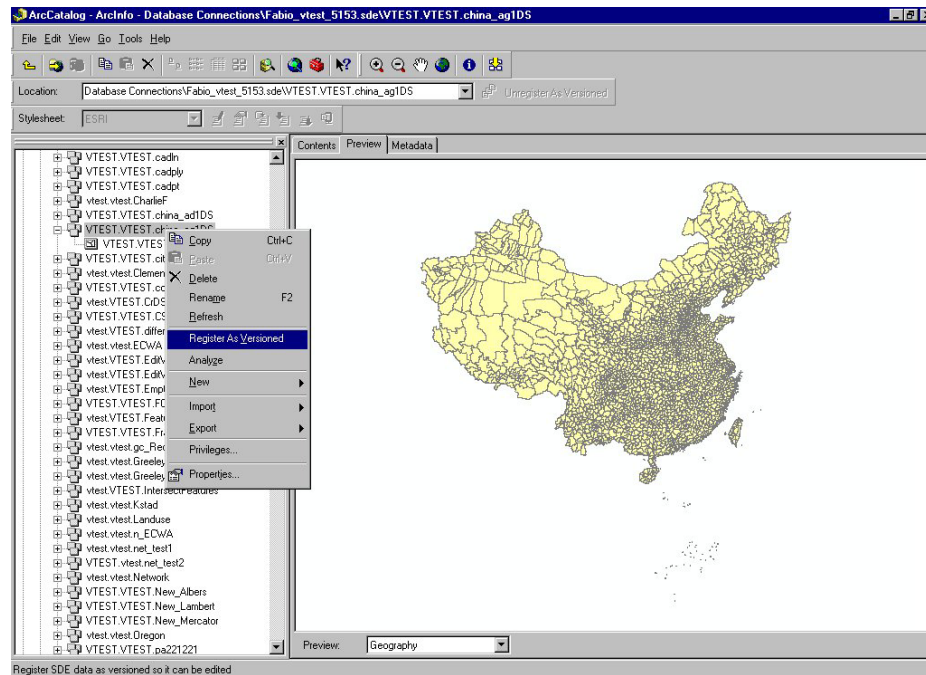
In the ArcToolbox Shapefile to Geodatabase wizard, we can see that a configuration keyword has been specified for the loading of the `hampton_streets` shapefile into the Geodatabase. Since the Geodatabase is maintained by an ArcSDE 8.3 service operating on a SQL Server database, we can store the resulting feature class using the parameters identified by the `Hampton_subdivision` `dbtune` keyword.

The shapefile CASNBRST.shp is converted to a feature class vtest.CASNBRST using ArcToolbox.



Versioning your data

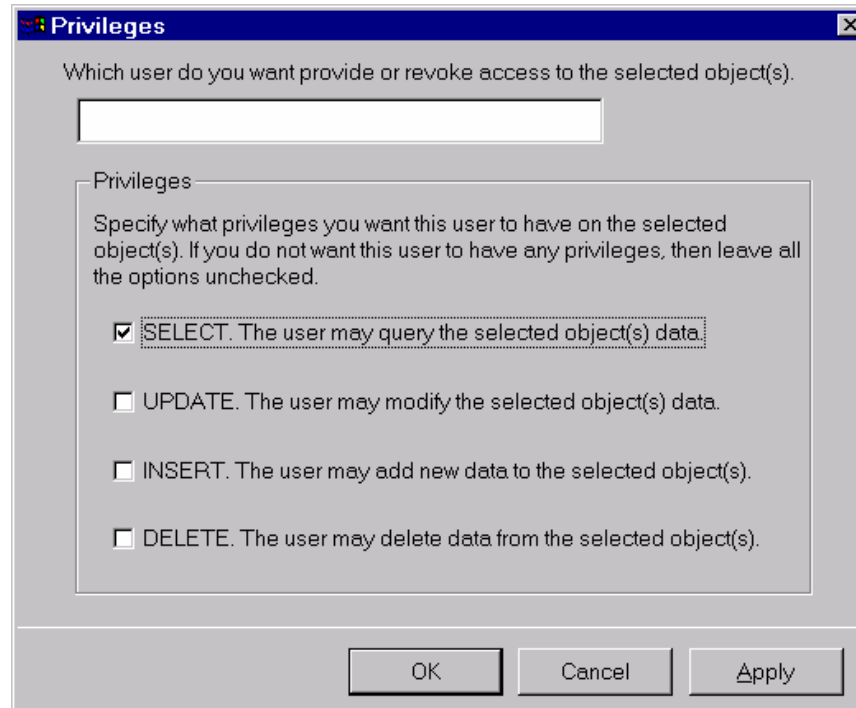
ArcCatalog also provides a means for registering data as versioned. Simply right-click the feature class to be registered as versioned and select the Register As Versioned context menu item. You can register standalone feature classes or all feature classes of a feature dataset.



A feature class is registered as multiversioned from within ArcCatalog.

Permissions

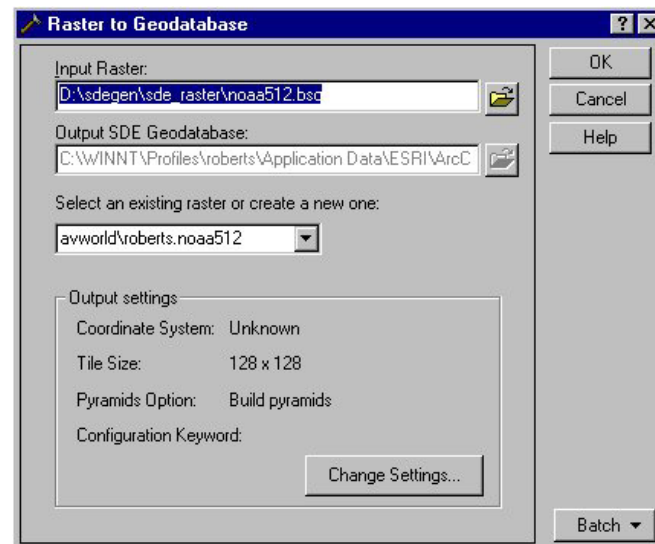
Using ArcCatalog, right-click on the data object class and click on the Privileges context menu. From the Privileges context menu assign privileges specifying the username and the privilege you wish to grant to or revoke from a particular user.



The ArcCatalog Privileges menu allows the owner of an object class, such as a feature dataset, feature class, or table, to assign privileges to other users or roles.

Loading a raster into ArcSDE with ArcCatalog

Using ArcCatalog, right-click on the database connection, point to Import, and click on Raster to Geodatabase. Navigate to the raster file to import. Click Change Settings if you want to change the coordinate reference system, tile size, pyramids option, or configuration keyword. Click OK to import the raster file into the SQL Server database.



Defining ArcSDE views

You can create spatial views of your data much the same way you create tabular views. Views are stored queries or virtual tables. The select statement is stored in the database instead of the table. SQL Server 2000 allows users to create unique clustered indexes on views and partition views across multiple servers.

Indexed views are stored in the database in the same manner as a table, while nonindexed views store only the SQL statements composing them.

With ArcSDE, it is possible to define views between two feature classes or between a feature class and a table or create more complex views containing subqueries or that span databases. ArcSDE views are created with the `sddtable -o create_view` command. When you create an ArcSDE view, three views are created, one of the business table, feature table, and spatial index.

Creating an ArcSDE view

Use the `sddtable -o create_view` command to create a view and the `sddtable -o delete` command to remove one. The syntax of `sddtable -o create` is:

```
sddtable -o create_view
  -T <view_name>
  -t <table1,table2...tablen>

  -c <table_coll,table_col2...table_coln>

  [-a <view_coll,view_col2...view_coln>]
  [-w "where_clause"]
  [-i <service>]
  [-s <server_name>]
  [-D <database>]
  -u <DB_User_name> [-p <DB_User_password>] [-N] [-q]
```

The syntax of `sddtable -o delete` is:

```
sddtable -o delete      -t <table>
  [-i <service>] [-s <server_name>] [-D <database>]
  -u <DB_User_name> [-p <DB_User_password>] [-N] [-q]
```

You must list the columns (-c) you want in the view as well as a where clause (-w) within the `sddtable` command. Be sure to include the business table's spatial column to make the view spatial. You must have 'create view' permissions to execute this command.

This example creates a view of all U.S. counties with a 1990 population per square mile greater than 50. The -s option specifies a server.

```
sddtable -o create_view -T CountyPopView -t us_counties -c
shape,name,state_name -w "pop90_sqmi > 50" -s bigman -u beetle -p bug
```

This example uses a subquery to select the state whose capital city has the largest male population.

```
sddtable -o create_view -T ManView -t us_states -c
us_states.shape,us_states.state_name --w "us_states.state_name = (select
us_capitals.state_name from us_capitals where p_male = (select max(p_male)
from us_capitals))" -u beetle -p bug -s bigman
```

Cross database views

It is possible to create views between tables and sde feature classes that do not reside in the same database. When you do this, you must correctly qualify the tables involved in your queries and use the -D (database) switch to identify which database contains the data on which to create a view.

This is illustrated in the following example, where there is a feature class in database "a." and a standalone, unregistered table in database "b.", and you want to make an sde view from these two entities in database "a."

1. Fully qualify the standalone table and its columns in your sdetable -o create_view statement. Or you can use table aliases. You must enclose table names in quotes if you do.
2. Qualify only to owner.tablename for the sde feature class.
3. Use the -D <database> switch for the "a" database.
4. Run the command as the owner of the sde layer.

The SQL statements to execute this would be:

```
CREATE VIEW crossDBView
AS SELECT  us_statesview.feature, us_statesview.state_name,
vtest.vtest.bob.sub_region
FROM  us_statesview, vtest.vtest.bob
WHERE  us_statesview.state_name=vtest.vtest.bob.state_name
```

So your command line would look like this:

```
C:\>sdetable -o create_view -T bigbob -t
vtest.us_statesview,vtest.vtest.bob
-c
us_statesview.feature,us_statesview.state_name,vtest.vtest.bob.sub_region
-w us_statesview.state_name=vtest.vtest.bob.state_name -i sql82k -D sde -u
vtest -p go
```

Explanation:

-T bigbob = name of the view

-t vtest.us_statesview,vtest.vtest.bob = us_statesview resides in the
'sde'
database, bob in 'vtest' database

-c

us_statesview.feature,us_statesview.state_name,vtest.vtest.bob.sub_region
= column list fully qualified for stand alone table.

-w us_statesview.state_name=vtest.vtest.bob.state_name = where clause,
fully
qualified for stand alone table.

-D sde = database context

Example using table aliases:

```
C:\sde81\etc>sdetable -o create_view -T crossdbview -t
vtest.us_statesview,"vtest.vtest.bob x" -c
us_statesview.feature,us_statesview.state_name,x.sub_region -w
"us_statesview.state_name=x.state_name and x.pop1996<1000000" -i sql82k -D
sde
-u vtest -p go
```

Creating views across linked servers

This is similar to cross database views with the following exceptions. You are forced to use aliases in your column lists because SQL Server will not allow four- or one-part qualified names in a column list. You must also define a linked server first and make

sure your DTC service is running. See the SQL Server Books Online for the exact usage of 'sp_addlinkedserver.'

1. Define a linked server in SQL Server Query Analyzer. This example was executed by a sysadmin user in SQL Server 2000.

```
sp_addlinkedserver
@server='foo',
@srvproduct='',
@provider='SQLOLEDB',
@datasrc='fabio\sql2k',
@location='',
@provstr = 'DRIVER={SQL
Server};SERVER=fabio\sql2k;UID=vtest;PWD=go;database=vtest',
@catalog='vtest'
```

2. Start the Distributed Transaction Coordinator (MSDTC). Open the Services panel and scroll down to the MSDTC service. If it is not running, start it.
3. Run the sdetable command with aliased table names:

```
C:\sde81\etc>sdetable -o create_view -T viewRemoteData -t
vtest.us_statesview,"foo.vtest.vtest.bob x" -c us_statesview.feature,
us_statesview.state_name,x.sub_region -w "us_statesview.state_name =
x.state_name and x.pop1996<1000000" -i sql82k -D sde -u vtest -p go
```

The SQL statements to execute this would be:

```
CREATE VIEW viewRemoteData AS SELECT
us_statesview.feature, us_statesview.state_name, x.sub_region
FROM vtest.us_statesview, foo.vtest.vtest.bob x
WHERE us_statesview.state_name=x.state_name
```

Creating indexed views (SQL Server 2000)

Indexed views are a new feature of SQL Server 2000 and are advantageous when your view definition employs a complex join. They are a good choice on static data as they are schemabound to their source. Also, they are stored in the same way as a regular table with a clustered index. Before you try this, make sure you are familiar with the SQL Server Books Online documentation covering indexed views.

Creating an indexed view in SQL Server 2000 requires extra work and has many conditions. These conditions can be summarized as follows:

- ⌘ The view must be created with Schemabinding.
- ⌘ The view must be created with QUOTED_IDENTIFIER on.
- ⌘ The view cannot reference other views.
- ⌘ Tables referenced in views must be two-part.
- ⌘ You cannot reference all columns of a table with *.

There are additional rules and conditions. Refer to the SQL Server Books Online under the topic 'Creating an Indexed View' for more information.

As ArcSDE does not create schema bound views by default, follow this process:

1. Create an ArcSDE view with sdetable -o create_view.
2. Open the SQL Server Query Analyzer and drop the business table view with drop view <view name>. Before you do this, make sure you copy the view definition.

3. Re-create the business table view with the create view statement, but make sure that you use the WITH SCHEMABINDING option and qualify your view name and table name with the owner.<table or view name> convention.
4. Create a unique clustered index on the view with the create unique clustered index statement.

This is an example that creates an indexed view. This example was done in SQL Server 2000 using a SQL Server authenticated user.

1. Create the view with sdetable -o create_view:

```
C:\sde81\etc>sdetable -o create_view -T vwWithIndx -t cagis_cent -c  
parcelid,object__id,x_coord,y_coord,book,page,parcel,shp -w  
"parcel>0393" -u vtest -p go -i sql82k -s fabio -D vtest
```

2. Query Analyzer—copy the view statement and then drop the view of the business table.

```
drop view vwWithIndx
```

3. Re-create the view. Make sure you qualify the view and table names with the owner name and use the WITH SCHEMABINDING statement.

```
CREATE VIEW vtest.vwWithIndx  
WITH SCHEMABINDING  
AS SELECT parcelid, object__id, x_coord, y_coord, book, page, parcel,  
shp  
FROM vtest.cagis_cent  
WHERE parcel>0393
```

4. Create a unique clustered index on this view:

```
create unique clustered index shpIdxView on vtest.vwWithIndx (shp)
```

Visualizing data with sdegroupp

If you have very large datasets, you'll want to discourage your users from drawing these at full extent. There is no use in drawing a million polygons at full extent, as you can see no detail. However, your users may want to visualize what the overall shape or extent of their data looks like. The command sdegroupp takes the existing shapes from a feature class and groups them into multipart shapes. While this may not be useful to analysis, it does allow rapid drawing of large amounts of data because it reduces the number of records. Polygons are generalized into polylines so as to not destroy topological relationships.

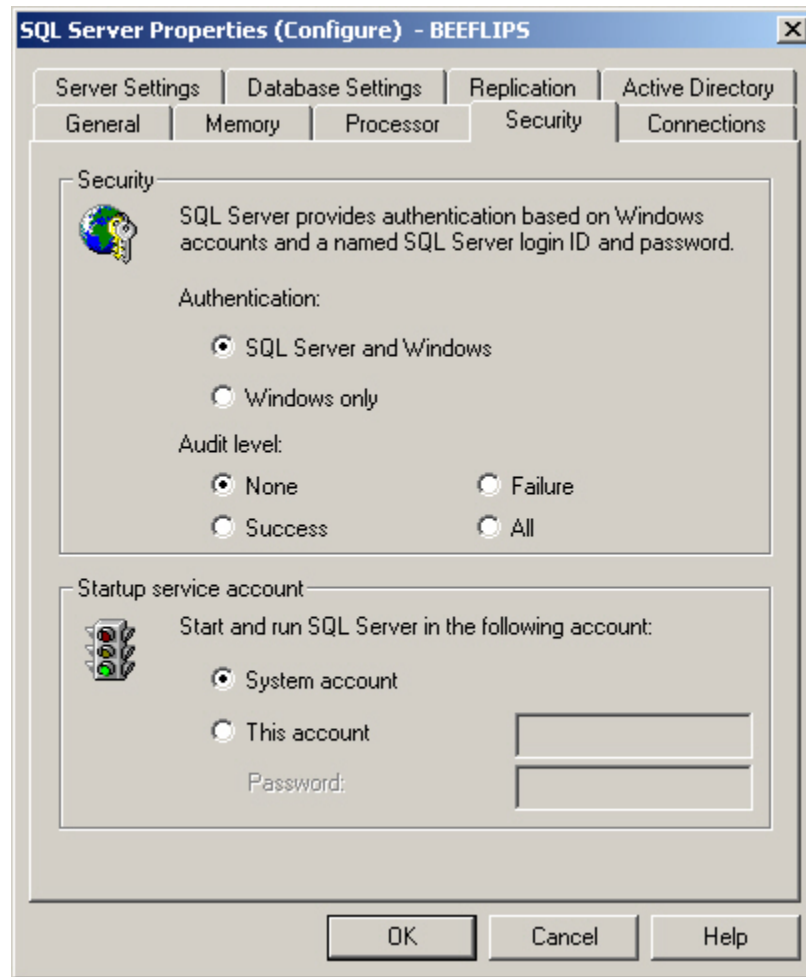
CHAPTER 5

Connecting to SQL Server

You can operate ArcSDE under three different configurations: three tiered with your clients connecting to a middle tier (giomgr) service; two tiered with your clients making direct connections to the database; and mixed, where your server accepts both direct and three-tier connections. This chapter covers how to make your initial connection to ArcSDE under a three-tiered configuration, how to connect in a two-tier configuration, and how to operate only in a two-tier environment without installing ArcSDE.

Making your first connection

VERY IMPORTANT NOTE: By default, SQL Server 2000 uses Windows authentication. In order to create the sde database and user, you must change this setting to include SQL Server authentication by right-clicking the server name in the Enterprise Manager and selecting Properties Security tab. Select, under Authentication, SQL Server and Windows. Then click OK to accept this setting. This is illustrated in the following panel. This procedure should be performed before installation of the ArcSDE software or creation of the sde database and sde user. After confirming that this option is correct, proceed with this manual.



If you have your ArcSDE service started, then you have already successfully made a connection to the data store (as the sde user). To successfully connect to an ArcSDE data store as a non-sde user, these conditions must be true:

1. The connecting user has been added as a user to the sde database. The connecting user cannot use a reserved word for a name. Examples of reserved words are max, min, national, and select. Refer to the SQL Server Books Online, under “reserved keywords,” for more information.
2. The connecting user has at least create table permissions in the sde database.
3. If the connecting user will own spatial data—that is, load spatial data with the ArcGIS data loaders such as ArcToolbox or the admin tools such as shp2sde or sdeimport—that user must have create table and create procedure permissions in the database that will store the loaded data.
4. You must explicitly connect to a database to perform DDL statements (create table, alter table, drop table). For example, if you want to load data into a database named county_hiways, you must connect to it by specifying a database name in your connection. ArcSDE does not use a login’s default database!

Example: Connecting to ArcSDE from ArcCatalog

Open ArcCatalog. Navigate to the **Database Connections** node. Expand it and select ‘Add Database Connection.’ Double-click this to bring up the connection properties dialog box. Under server, enter the name of the ArcSDE server. For Service, input either your service instance name or your tcp/ip port (e.g., 5151). For Database, leave blank to connect to the sde database or input another database name. Input a user name and password for the remaining text boxes.

Figure 5.1: Connecting to a server through ArcCatalog

Spatial Database Connection

Server: offwidth

Service: 5152

Database:
 (If supported by your DBMS)

Account

User Name: gdb

Password: xxxx

☒ Save Name/Password Test Connection

Version

☒ Save Version

sde.DEFAULT Change...

OK Cancel

Example: Connecting to ArcSDE through ArcObjects™

To use this sample, you must have the esricore.tlb referenced in your MS Visual Basic development project.

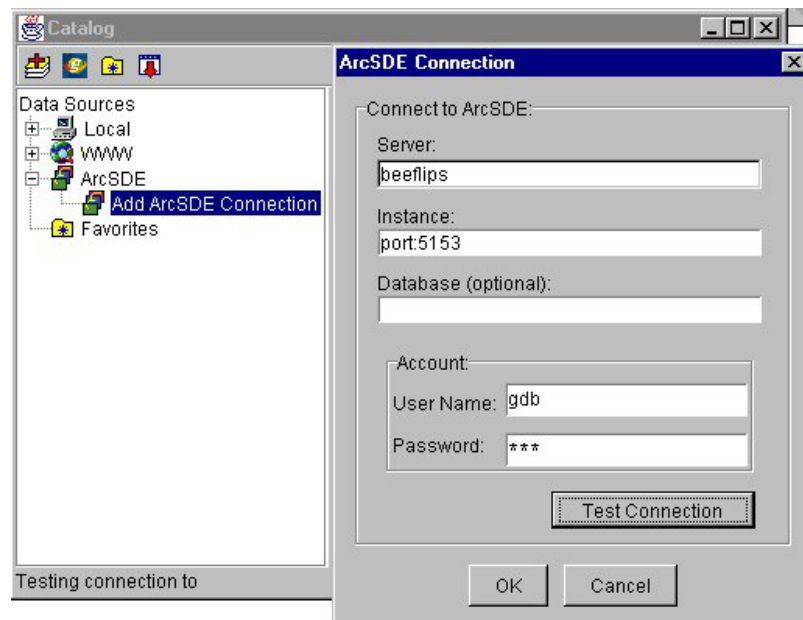
```
Dim pPropSet as IPropertySet
Dim pWorkspaceFactory as IWorkspaceFactory
Dim pWorkspace as IWorkspace

set pPropSet = new PropertySet
set pWorkspaceFactory = new SDEWorkspaceFactory
with pPropSet
    .setProperty "Server" = "vogon"
    .setProperty "Instance" = "5151"
    .setProperty "Database" = "contours"
    .setProperty "User" = "arthur"
    .setProperty "Password"="dent"
    .setProperty "version" = "sde.DEFAULT"
end with
set pWorkspace=pWorkspaceFactory.Open(pPropSet)
if pWorkspace is nothing then
    msgbox "Connection Failed"
else
    msgbox "Success"
end if
```

Example: Connecting to ArcSDE through ArcExplorer™

Open the Catalog, expand the ArcSDE servers node, then double-click the Add ArcSDE Connection entry. In the ArcSDE Connection panel, fill out your server, instance, database (if you leave this blank, you'll connect to the 'sde' database), user name, and password arguments, then click "Test Connection." If the connection is successful, (grayed out), then click OK.

Figure 5.2: Connecting to ArcSDE from ArcExplorer



Using the ArcSDE Direct Connect Driver

Beginning with ArcSDE 8.1, ArcSDE client applications can connect to either an ArcSDE service or directly to an MS SQL Server instance. The direct connection capability, added during the ArcSDE 8.1 release, was achieved by embedding the ArcSDE application server technology into the ArcSDE client software. Any application programmed with the functionality of the ArcSDE 8.1 C API library can connect directly to an MS SQL Server instance.

ArcSDE client applications that connect to an ArcSDE service send spatial requests to the service. The ArcSDE service converts the spatial requests into SQL statements and submits them to the SQL Server instance. Upon receiving results from the server, the ArcSDE service converts the result into spatial data recognizable to the ArcSDE client. At ArcSDE 8.1, the functionality of the ArcSDE service has been linked into the ArcSDE C API, allowing ArcSDE client applications to connect directly to MS SQL Server instances. Spatial requests are converted to SQL statements by the ArcSDE client applications instead of the ArcSDE service.

Before you can connect to an ArcSDE service, you must install and configure the ArcSDE product. The sde user must be created and the sde setup program (sdesetupmssql.exe) must be run to create and populate the ArcSDE system tables and required stored procedures. In addition, ArcSDE must be able to reference an ArcSdeServer license from a license manager accessible through the network. Since a direct connection to an MS SQL Server instance does not require the presence of an ArcSDE service, the ArcSDE product does not need to be installed. You are, however, required to run the ArcSDE setup program, which creates the necessary ArcSDE system tables in the sde user's schema. Also, to obtain a read-write connection to the MS SQL Server database, an ArcSDE client application must reference a valid ArcSdeServer license. If an ArcSdeServer license cannot be referenced, access is restricted to read-only.

The ArcSDE setup program is located in the bin directory of all ESRI products capable of connecting to ArcSDE. If you do not already have one, an ArcSdeServer license can be obtained from ESRI Customer Support. To connect directly to an MS SQL Server instance, you must have Microsoft Data Access Components 2.5 installed on your client machine. This MDAC is installed by ArcGIS, so you may already have it. The sections that follow provide details on how to set up the client and server machines to perform a direct connection to an MS SQL Server instance. For more information regarding an ArcSDE service, refer to the *ArcSDE Installation Guide* for installation instructions and the *Managing ArcSDE Services* for configuration and connection instructions.

All ArcSDE client applications, such as ArcMap or ArcCatalog, function the same way regardless of whether a user connects to an ArcSDE service or directly to an MS SQL Server instance. The only difference occurs during the entry of the connection information.

Setting up your server without installing ArcSDE

ArcSDE requires an 'sde' database and at least an 'sde' user with create table, view, and stored procedure permissions. The two most common ways of doing this are with the Enterprise Manager for SQL Server 7 or SQL Server 2000 or with a t-sql script. After creating your sde database and user, you must run the sdesetupmssql in the \bin directory where you have the gsrvsrql81.dll; the sdesqlsvr81.dll; and three client libs sg81.dll, pe81.dll, and sde81.dll. Finally, set either SDEHOME or ARCHOME to the directory containing the bin folder holding these files.

1. Create an ArcSDE database and user

Either run this script in the SQL Server Query Analyzer or see the screenshots from the SQL Server Enterprise Manager below:

NOTE: If you run this script, it will be necessary to change your paths where indicated.

```
/*
Script: setupSDEdb.sql
purpose:
1. Creates sde database if it doesn't exist
2. Creates sde and gdb logins if they don't exist
3. adds logins to sde database if they are not there
4. grants create table view and procedure to users in sde.

NOTE: Change the paths in the filename statements below
These files will be written to your disk
*/

use master
go

if exists (select name from sysdatabases where name = 'sde')
begin
    print 'SDE database already exists, attempting to create login and user'
end
else
begin
    CREATE DATABASE sde
    on
    ( name='sde',
    --change paths here
    --to a valid location on
    --your disk
    filename='d:\sdedb\sql\sde.mdf',
    size=30,
    maxsize=UNLIMITED,
    filegrowth=30 )
    LOG ON
    ( name='sdelog',
    --change paths here
    --to a valid location on
    --your disk
    filename='d:\sdedb\sql\sdelog.ldf',
    size=10mb,
    maxsize=UNLIMITED,
    filegrowth=10mb)
end

/*
now attempt to create logins and users
if they don't exist
*/

if exists (select name from syslogins where name = 'sde')
begin
    print 'sde login exists, attempting to create test user'
end
else
begin
    --create sde login
    --password is 'go'
    --default db = sde
    exec sp_addlogin 'sde', 'go', 'sde'
end

--create a test user
if exists (select name from syslogins where name = 'gdb')
begin
    print 'gdb test login already exists'
end
else
begin
    --create a test login whose name = gdb
    --and password = gdb
    exec sp_addlogin 'gdb','gdb','sde'
```

```

end

/*
switch db context to sde
check for users in sde
add if necessary
grant create table, view and procedure to
both users
*/
go
use sde
go

--sde login
if exists (select name from sysusers where name = 'sde')
begin
    print 'sde user already added to sde database'
end
else
begin
    exec sp_adduser sde
end

--gdb login
if exists (select name from sysusers where name = 'gdb')
begin
    print 'gdb user already added to sde database'
end
else
begin
    exec sp_adduser gdb
end
go

grant create table to sde
grant create table to gdb
grant create view to sde
grant create view to gdb
grant create procedure to sde
grant create function to sde --SQL SERVER 2K ONLY!!
grant create procedure to gdb

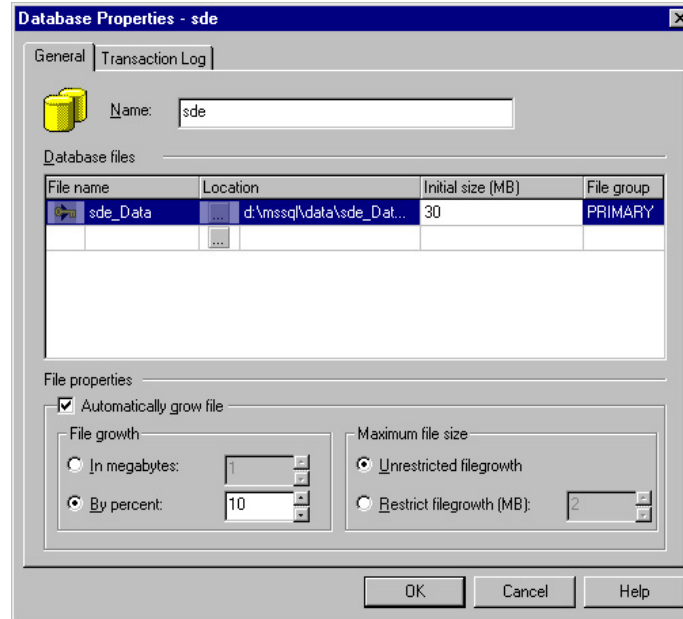
/*
make sure you run sdesetupmssql -o install
from either %sdehome%\bin or %archome%\bin
*/

```

2. Create the SDE database and user with the SQL Server Enterprise Manager

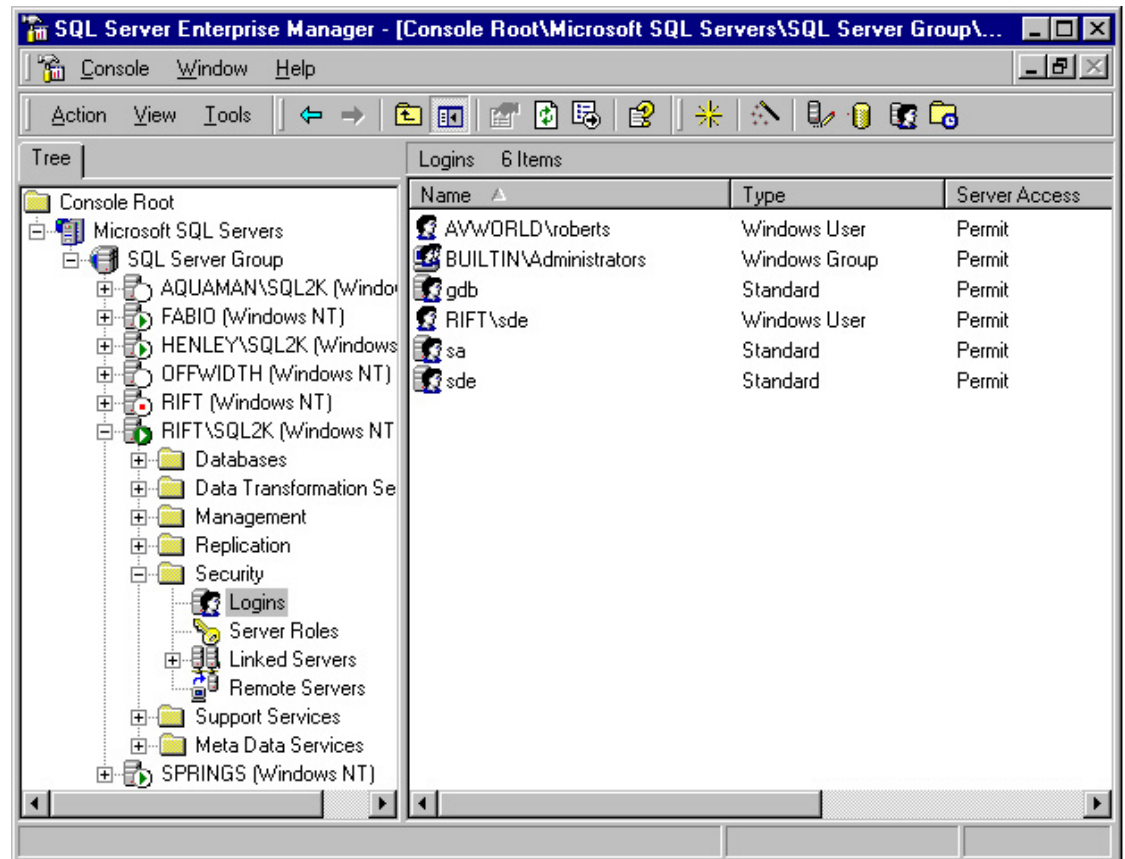
Create the sde database

Open Enterprise Manager, expand your “Microsoft SQL Servers” node, expand your server, right-click Databases, and select New Database.



You can take the defaults here or fill out the initial size, location, etc. Do the same for the transaction log. Click OK and the sde database will be created.

Create the sde login and user: From within your server's node, expand the security node and select logins.

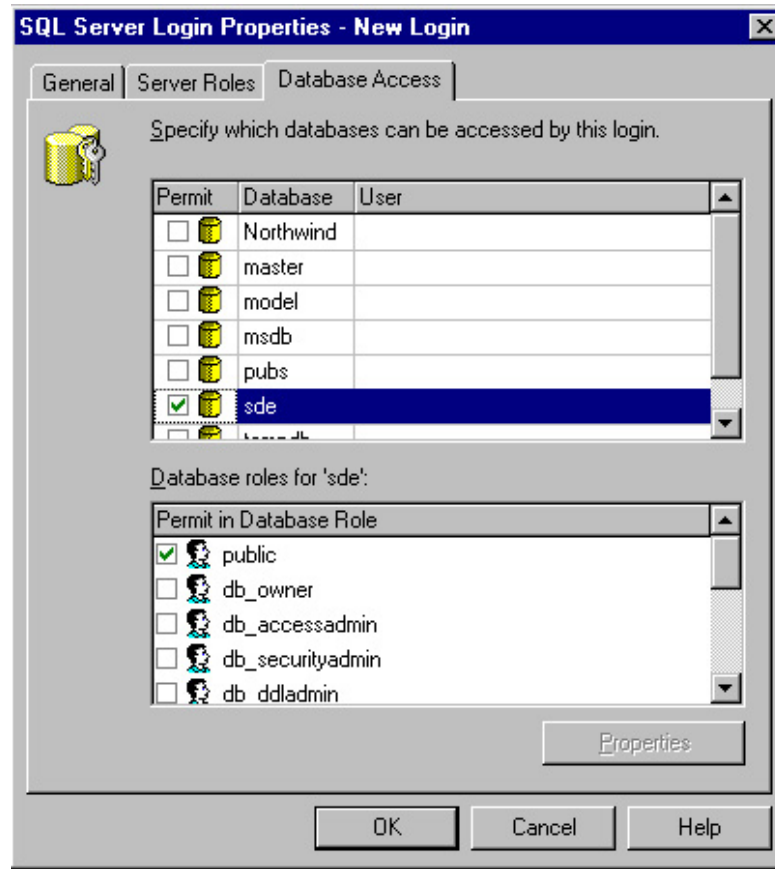


Now right-click in the right-hand panel. Select New Login.

The screenshot shows the 'SQL Server Login Properties - New Login' dialog box. It has three tabs: 'General', 'Server Roles', and 'Database Access'. The 'General' tab is active. It contains a 'Name' field with 'sde' entered. Below it is the 'Authentication' section with two radio buttons: 'Windows Authentication' (unselected) and 'SQL Server Authentication' (selected). Under 'Windows Authentication' are 'Domain' and 'Security access' (Grant/Deny) options. Under 'SQL Server Authentication' is a 'Password' field with 'xx' entered. At the bottom is the 'Defaults' section with 'Database' (sde) and 'Language' (<Default>) dropdowns. 'OK', 'Cancel', and 'Help' buttons are at the bottom right.

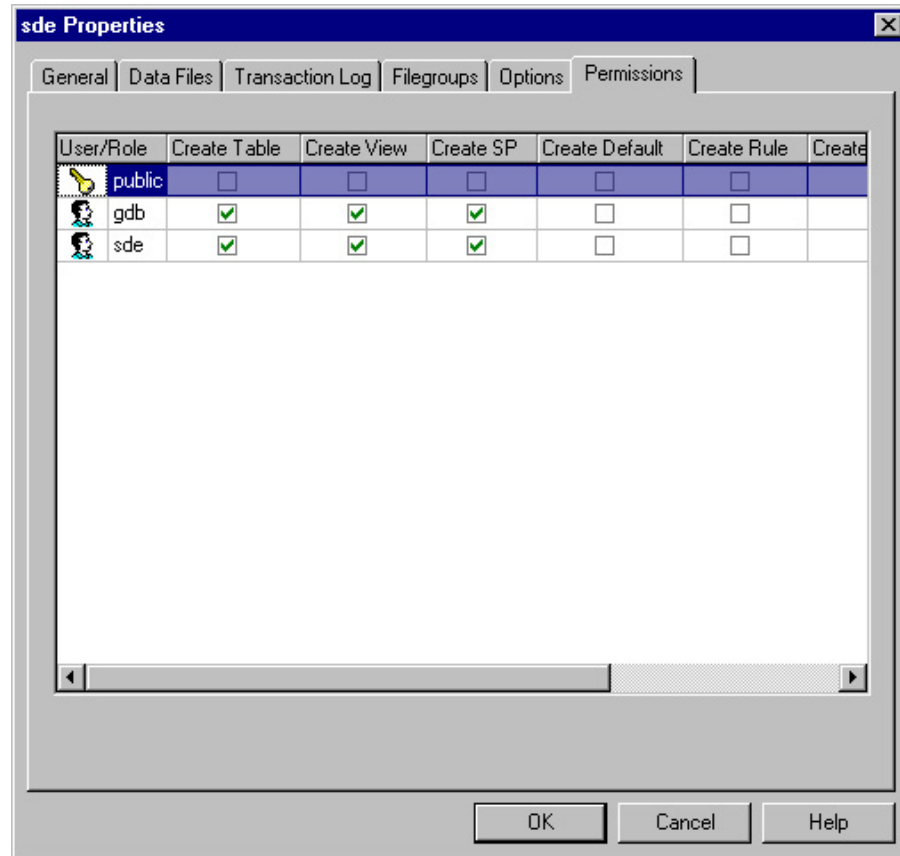
Fill in the login name property and password.

The default database should be 'sde.' Now click the 'Database Access' tab and click the sde database check box.



Now grant this new user create table, view and stored procedure in the 'sde' database. Go back to the database's node, right-click the 'sde' database, and select Properties. When the Properties dialog appears, click the Permissions tab.

Check the create table, view, and sp (stored procedure) check boxes. If you are using SQL Server 2000, you must also grant create function to sde. Click OK.



RUN sdesetupmssql

You are almost done. Now from a command prompt, navigate to the directory that contains `gsrvrsql81.dll`, `sdesqlsvr81.dll`, `sg81.dll`, `sde81.dll`, and `pe81.dll` with a `cd (chdir)` command. There should also be an executable in this directory called `sdesetupmssql.exe`. In this example, these executables are located in the ArcGIS installation folder. This executable creates the sde and Geodatabase metadata. The syntax for this command is listed below.

```
C:\ArcGIS\arcexe81\Bin>sdesetupmssql
Usage for sdesetupmssql:
-?
-h
-o upgrade [-H <sde_directory>] [-u <DB_Admin_user>] [-p
<DB_Admin_password>]
[-s datasource] [-N] [-q]
-o list [-H <sde_directory>] [-u <DB_Admin_user>] [-p <DB_Admin_password>]
[-s datasource] [-q]
-o install [-H <sde_directory>] [-u <DB_Admin_user>] [-p
<DB_Admin_password>]
[-s datasource] [-N] [-q]
```

In this example below, `sdesetupmssql` will be run against a data source called “rift/sql2k,” which is a SQL Server 2000 instance name.

```
C:\ArcGIS\arcexe81\Bin>sdesetupmssql -o install -u sde -p go -s rift/sql2k
-H c:\ArcGIS\arcexe81
```

The user should confirm that the `-H` switch is entered properly.

How do you know what your data source is? The data source will almost always be the name of your SQL Server host. However, if you have both SQL Server 7 and SQL Server 2000 installed, the SQL Server 2000 instance will have a different name that you assigned to it when you did the installation. If this is the case, look in the Enterprise Manager. You should see both your instances listed here.

Making a direct connection

Whether or not you have an ArcSDE service running or have installed the ArcSDE software, you are ready to make a direct connection. Here's what you must do:

You must have sdehome set and pointing to the folder that contains the bin directory holding the gsrvsrql81.dll; the sdesqlsvr81.dll; and the client libraries sg81.dll, sde81.dll, and pe81.dll.

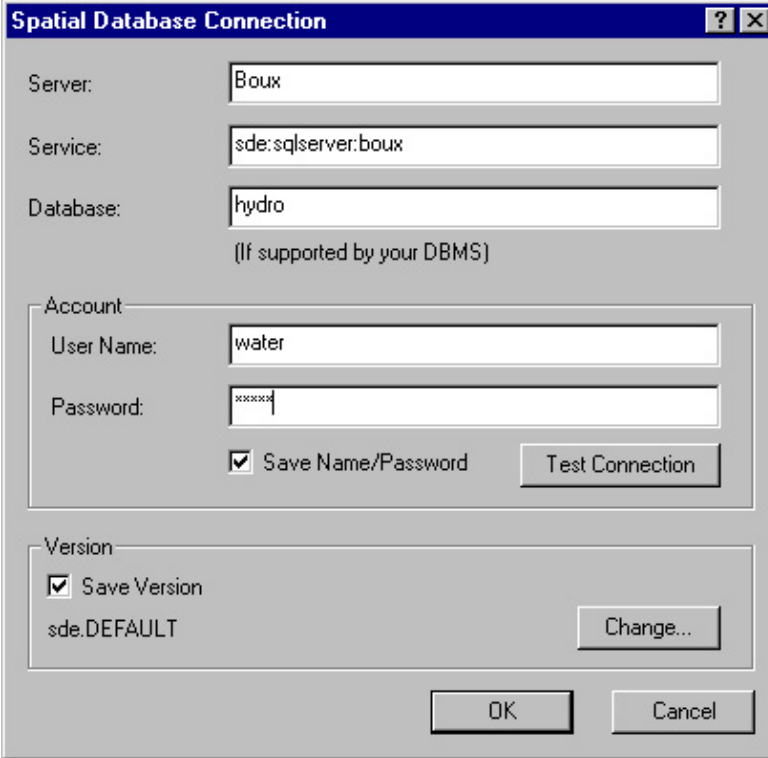
- € You can use substitute SDEHOME with ARCHOME.
- € Or set neither and let the search default to HKEY_LOCAL_MACHINE-SOFTWARE-ESRI-ArcInfo-Desktop-8.0-SourceDir - BIN.

For example: Presume you are attempting a direct connection from ArcCatalog. You have installed ArcGIS into c:\arcgis\arcexe81. The bin folder beneath the arcexe81 directory holds the gsrvsrql81.dll, the sdesqlsvr81.dll, and the three client libs. You can set SDEHOME to c:\arcgis\arcexe81 and attempt your direct connection.

If you have one setup, try to connect. To connect:

- Server name = ArcSDE Server (or any string)
- **Instance = sde:sqlserver:<data source>**
- Database = a database you want to explicitly connect to
- User = any valid DBMS/ArcSDE user
- Password = any valid DBMS/ArcSDE user's password

You must have at least MDAC 2.5 installed. How can you tell? Check the file version of msado15.dll in \program files\common files\system\ado. If its version is 2.5 or greater, you are OK. You can also use the component checker from Microsoft. This tool will allow you to remove MDAC as well. Beware of installing an MDAC onto Windows 2000; it may already have it installed.



The image shows a 'Spatial Database Connection' dialog box. It has a title bar with a question mark and a close button. The dialog is divided into several sections. The top section contains three text boxes: 'Server' with 'Boux', 'Service' with 'sde:sqlserver:boux', and 'Database' with 'hydro'. Below the 'Database' box is the text '(If supported by your DBMS)'. The middle section is titled 'Account' and contains 'User Name' with 'water' and 'Password' with 'xxxxxx'. There is a checkbox labeled 'Save Name/Password' which is checked, and a 'Test Connection' button. The bottom section is titled 'Version' and contains a checkbox labeled 'Save Version' which is checked, and a text box with 'sde.DEFAULT'. There is a 'Change...' button next to the version text box. At the bottom of the dialog are 'OK' and 'Cancel' buttons.

Spatial Database Connection

Server: Boux

Service: sde:sqlserver:boux

Database: hydro
(If supported by your DBMS)

Account

User Name: water

Password: xxxxxx

☒ Save Name/Password Test Connection

Version

☒ Save Version

sde.DEFAULT Change...

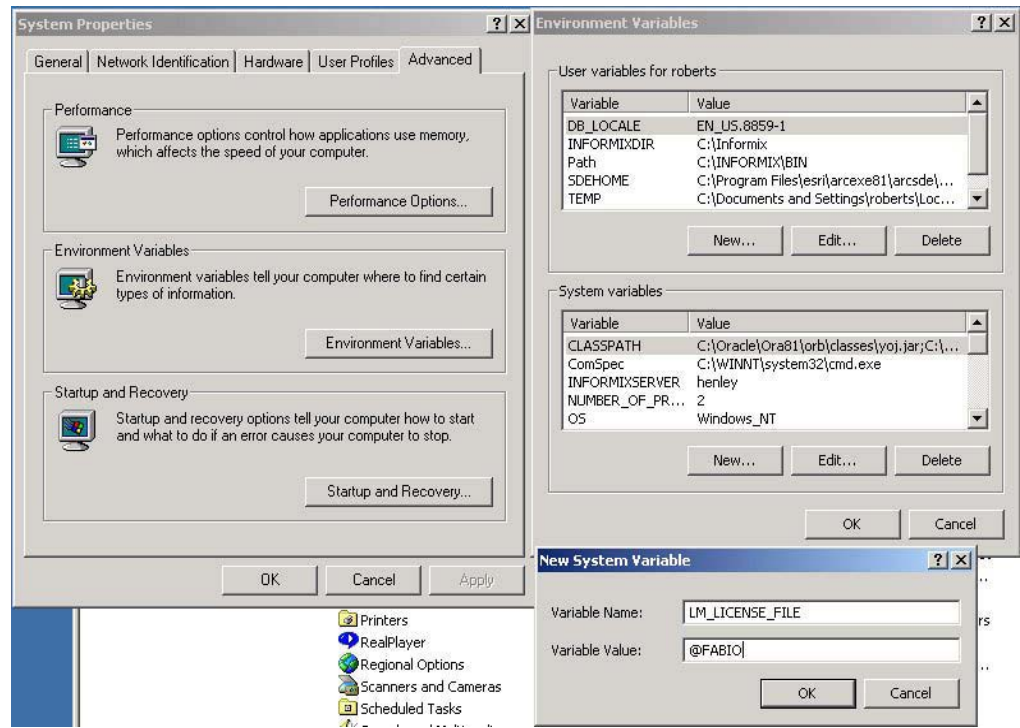
OK Cancel

Licensing Direct Connect clients

If you do not have a valid ArcSDE license service, you will still be able to connect to any ArcSDE data store. However, your connections will be read-only. If you do have a license server, you can set an environment variable, `LM_LICENSE_FILE`, to access those licenses. Make sure you set it with the host name prefixed with an @ sign. For example, setting this variable in DOS would look like this:

```
C:\sde>SET LM_LICENSE_FILE=@FABIO
```

Setting this variable in Windows 2000 would look like the following panel.



CHAPTER 6

National language support

Storing data in an ArcSDE SQL Server database using collation sets other than the SQL Server default, Latin1_General, requires some extra configuration on both the client and the server. This section provides guidelines for configuring both the SQL Server database and the ArcSDE client environment to enable the use of collation sets other than Latin1_General.

SQL Server database collation designator

The default collation designator for SQL Server 2000 databases is Latin1_General. The collation is selected when SQL Server 2000 is installed but can be changed later by running the rebuildm.exe utility. This utility is documented in the SQL Server Books Online. Running it will remove any existing databases. The default collation will support U.S. English and most Western European languages. Change this default only if you are not using a language from this group. Consult the SQL Server Books Online to determine the collation designator that is appropriate for your data. Also, select how data will be sorted. Sort order can be binary, case sensitive, accent sensitive, kana sensitive, or width sensitive. If the binary option is not selected, SQL Server 2000 follows sorting and comparison rules as defined in dictionaries for the associated language or alphabet. It is also possible to choose an alternative collation if backward compatibility with previous installations of SQL Server 7 or earlier is desired. The choice here will then depend on the collation choice made in the earlier instance of SQL Server. If the previous installation accepted the defaults, select Dictionary Order, Case Insensitive. This is backward compatible with the 1252 character set, which was the default in SQL Server 6.5 and SQL Server 7.0.

Collation is read from the master database. All ArcSDE spatial databases must share the same collation as the master and sde databases.

Setting the SDE_SQLCHARSET variable for Windows NT clients

Once the ArcSDE Sqlserver has been set up with the desired collation set, you can store and retrieve data from the database from any client as long as the client is using the same character encoding standard or codepage as the ArcSDE Sqlserver's host computer. If the ArcSDE client is running on a computer using a different character encoding, the client must set the SDE_SQLCHARSET variable, to be able to send and receive character data from the ArcSDE server.

For instance, let's assume the ArcSDE Sqlserver, running on Windows NT, is using the Shift JIS character encoding for Japanese characters. Any ArcSDE Client connecting to this Sqlserver database must also use Shift JIS character encoding. If there is a Unix

client using the EUC JP character encoding, the client must set the SDE_SQLCHARSET variable to be able to send and receive character data to and from the ArcSDE server.

For a C API or ArcSDE command utility application, the SDE_SQLCHARSET must be set as a system environment variable.

setenv SDE_SQLCHARSET eucjis

For Java API client applications, set the SDE_SQLCHARSET as a JVM system property.

java -DSDE_SQLCHARSET=eucjis [Java client program]

Setting the SDE_SQLCHARSET variable for Windows NT Clients

Exercise care in setting the SDE_SQLCHARSET on the Windows NT platform. In the Windows NT environment there are two different codepage environments, Windows American National Standard Institute (ANSI) and Original Equipment Manufacturer (OEM). That is, the sets of character encoding standards supported by Windows environment and DOS environment are different. Windows applications such as ArcInfo and ArcView run in the ANSI codepage environment. ArcSDE administration tools and C API applications invoked from the MS-DOS Command Prompt run in the OEM codepage environment. If you use both Windows applications and MS-DOS applications together then set the SDE_SQLCHARSET variable for MS-DOS applications by opening the MS-DOS Command Prompt and issuing a MS-DOS SET command. Initially verify that the data to be loaded is in either ANSI or OEM format. From the MS-DOS Command Prompt issue one of the following commands:

If the ArcSDE Windows NT client, using Baltic language as the local language, is running from Windows, set the codepage to cp1257 which is the ANSI codepage for Baltic languages.

set SDE_SQLCHARSET=cp1257

If the client is running from the MS-DOS prompt, set the codepage to cp775, which is the OEM codepage for Baltic language.

set SDE_SQLCHARSET=cp775

If the ArcSDE client is running from a UNIX platform, set the codepage to iso_13

setenv SDE_SQLCHARSET iso_13

Note that there should not be any blank spaces at the end of the command.

Character encoding standards supported by ArcSDE

Currently ArcSDE only supports conversions between the character encoding standards listed in the table below.

Character encoding	Description
big5	BIG5 16-bit Traditional Chinese
Euccns	EUC 32-bit Traditional Chinese
eucgb	CGB2312-80 16-bit Simplified Chinese ZHS16CGB231280
sjis	Shift-JIS 16-bit Japanese
euclj	EUC 24-bit Japanese
eucksc	KSC5601 16-bit Korean
iso_1	ISO 8859-1 West European
iso_2	ISO 8859-2 East European
iso_3	ISO 8859-3 South European
iso_4	ISO 8859-4 North and North-East European
iso_5	ISO 8859-5 Latin/Cyrillic
iso_6	ISO 8859-6 Latin/Arabic
iso_7	ISO 8859-7 Latin/Greek
iso_8	ISO 8859-8 Latin/Hebrew
iso_9	ISO 8859-9 West European & Turkish
iso_10	ISO 8859-10 North European
iso_13	ISO 8859-13 Baltic
iso_15	ISO 8859-15 West European
cp437	IBM-PC Code Page 437 8-bit American
cp850	IBM-PC Code Page 850 8-bit West European
cp851	IBM-PC Code Page 851 8-bit Greek/Latin
cp852	IBM-PC Code Page 852 8-bit East European
cp855	IBM-PC Code Page 855 8-bit Latin/Cyrillic
cp857	IBM-PC Code Page 857 8-bit Turkish
cp860	IBM-PC Code Page 860 8-bit West European
cp861	IBM-PC Code Page 861 8-bit Icelandic
cp862	
cp863	IBM-PC Code Page 863 8-bit Canadian French
cp864	
cp865	IBM-PC Code Page 865 8-bit Norwegian
cp866	IBM-PC Code Page 866 8-bit Latin/Cyrillic
cp869	IBM-PC Code Page 869 8-bit Greek/Latin
cp720	
cp737	IBM-PC Code Page 737 8-bit Greek/Latin
cp775	IBM-PC Code Page 775 8-bit Baltic
cp1250	MS Windows Code Page 1250 8-bit East European
cp1251	MS Windows Code Page 1251 8-bit Latin/Cyrillic
cp1252	MS Windows Code Page 1252 8-bit West European
cp1253	MS Windows Code Page 1253 8-bit Latin/Greek
cp1254	MS Windows Code Page 1254 8-bit Turkish
cp1255	MS Windows Code Page 1255 8-bit Latin/Hebrew
cp1256	MS Windows Code Page 1256 8-bit Latin/Arabic
cp1257	MS Windows Code Page 1257 8-bit Baltic
cp1258	MS Windows Code Page 1258 8-bit Vietnamese

CHAPTER 7

Standby servers and log shipping

The use of a standby server that can be brought online if the primary production server fails can avert the loss of data or allow users to continue working with the database. SQL Server 7 and SQL Server 2000 both support the use of standby servers. A standby server contains a copy of the databases on the primary server and can be brought online if the primary production server fails.

Implementing a standby server

Implementing a standby server involves three phases:

1. Creating the database and ongoing transaction log backups on the primary server.
2. Setting up and maintaining the standby server by backing up the database on the primary server and restoring it on the standby server.
3. Bringing the standby server online if the primary server fails.

If the use of a standby server is necessary, all users will have to log onto the standby server. User processes are not transferred automatically to the standby server. In addition, transactions are not maintained. It will be necessary to run `sp_change_users_login` once the switch to the standby server is made.

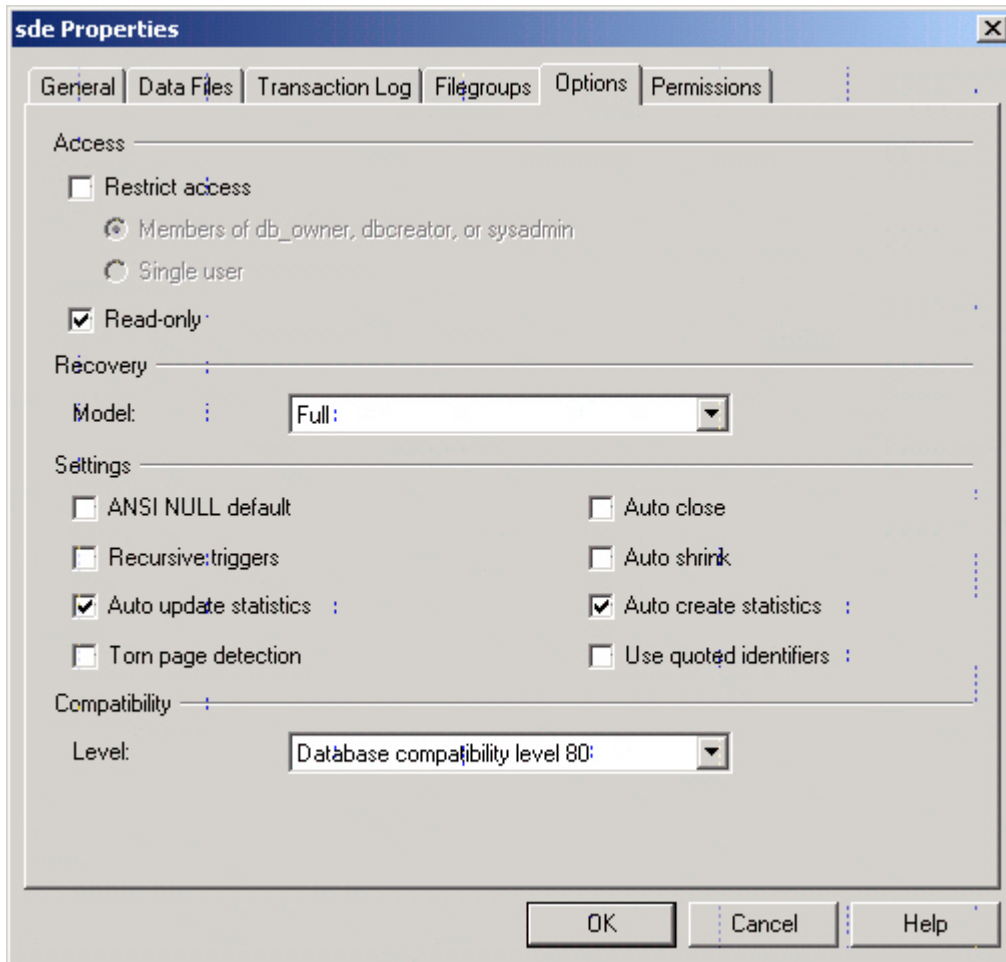
To test ArcSDE 8.1 in a Standby Server configuration, the following steps were implemented. For more information, see the instructions in SQL Server Books Online.

1. Create the full database backups on the Primary Server. Create a transaction log backup of each database that is duplicated. The frequency of these transaction log backups will depend on the amount of transaction changes that users make.
2. Restore the database backups onto the standby server that is in standby mode. Specify one undo file per database. Periodically apply the transaction logs from the primary server to the databases on the standby server. Specify the same undo file used in previous steps.

3. If the primary server should become unavailable, apply to the standby server any transaction logs that have not yet been applied to it. Recover the databases on the standby server.
4. The standby server is now available for users to make changes to the database. Configuration of the standby server can be done in either Enterprise Manager or Query Analyzer.

The following steps illustrate the creation of a standby server in Enterprise Manager.

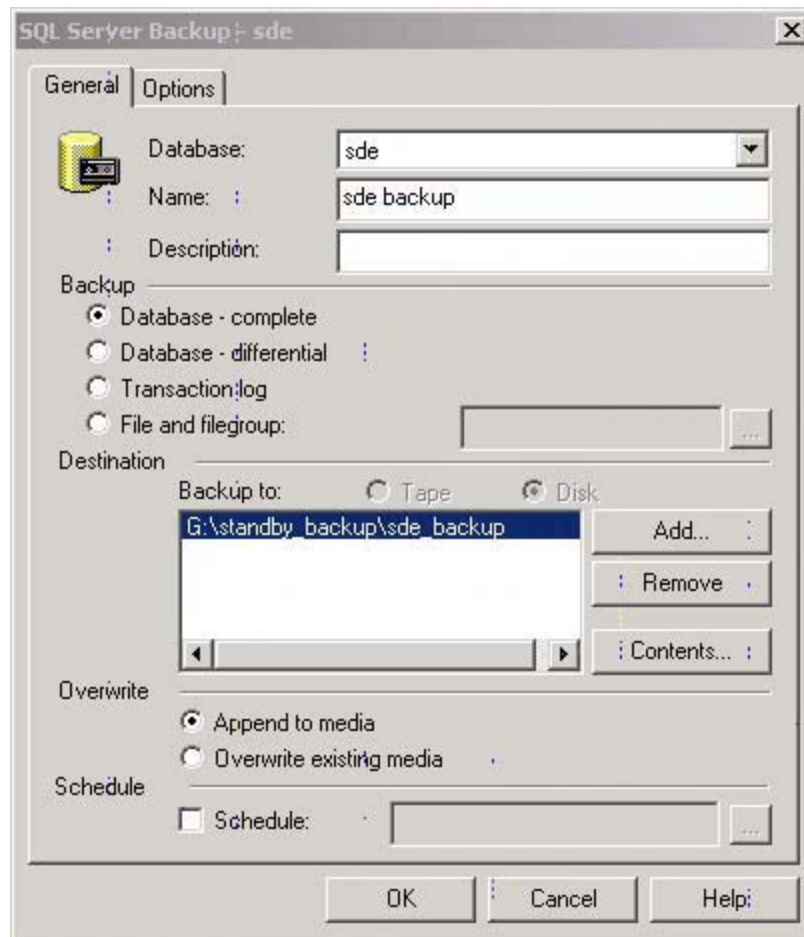
Confirm that the database is in Full Recovery Model.



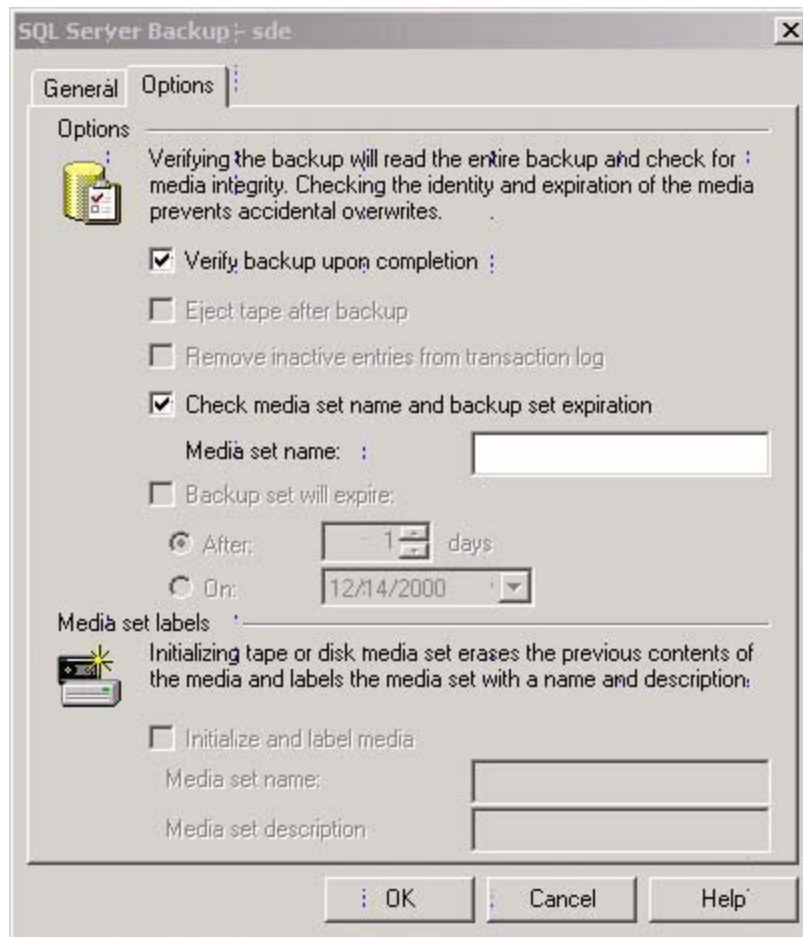
In order to implement transaction log backups, the database must be in Full Recovery Model. This is the default setting for a new database created with the Standard and Enterprise Editions of SQL Server 2000. To change the Recovery Model in Enterprise Manager, right-click on a database and choose Properties. Click the Options tab and change the recovery model to Full. Click OK to continue.

Backing Up Databases:

In Enterprise Manager, right-click on the database to backup, choose all tasks, Backup Database.

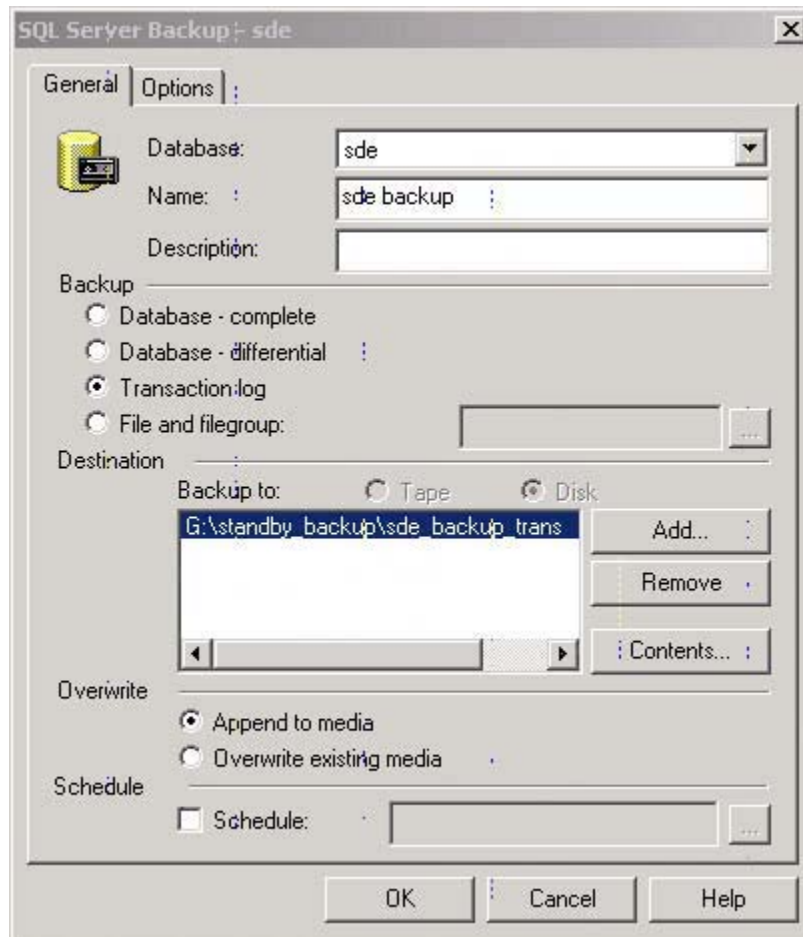


Choose Backup Database Complete and select the location in which to store the backup. Click on the Options tab.



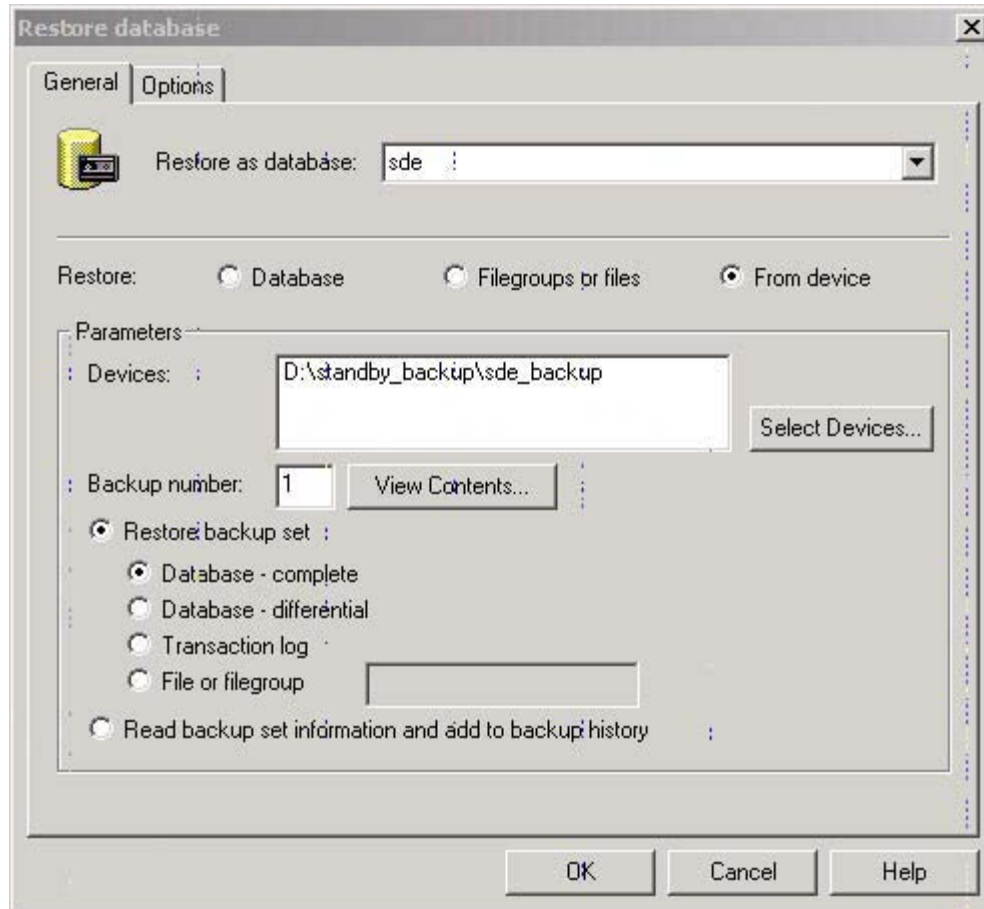
Click Verify backup upon completion and click OK to continue.

To back up the transaction log, right-click the database, choose all tasks, Backup Database.

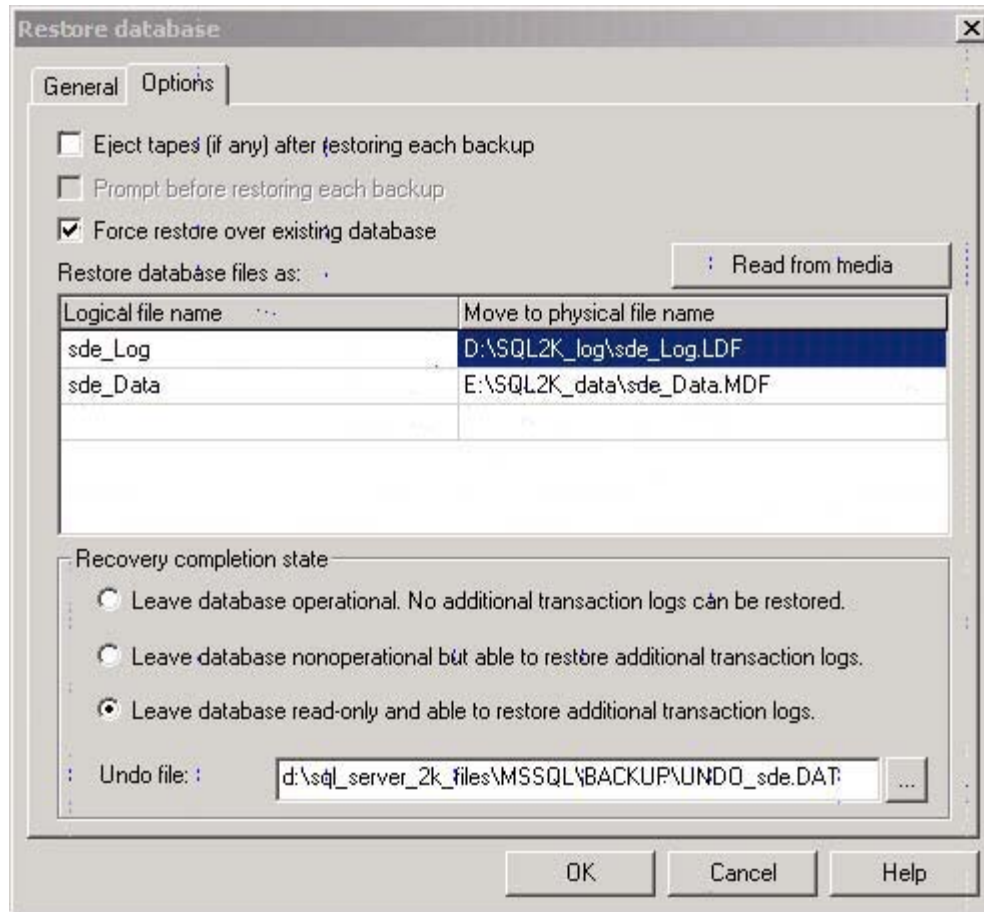


Click on Backup Transaction log. Choose the destination for the backup file. As before, click on the Options tab and click on Verify backup upon completion. Click on OK to continue.

To restore the database on the standby server, right-click on the database to restore, choose all tasks, restore database.



Choose Restore backup set Database - complete. Select Restore From device and then select the file to restore. Click on the Options tab.



In the Recovery completion state, select Leave database read-only and able to restore additional transaction logs. Choose the location and filename of the undo file. Click OK to continue. The database will be restored and will be in read-only mode. Continue this process for each transaction log to be restored on the standby server.

If the production server should become unavailable, apply all remaining transaction logs to the standby server. From Query Analyzer recover the standby server database without restoring. Execute the following statement with the name of the database to be recovered.

-- Restore database using WITH RECOVERY.

RESTORE DATABASE sde WITH RECOVERY

The standby server will now be in the same state as the production server from the last transaction log. All changes after the last backup of the transaction log will be lost.

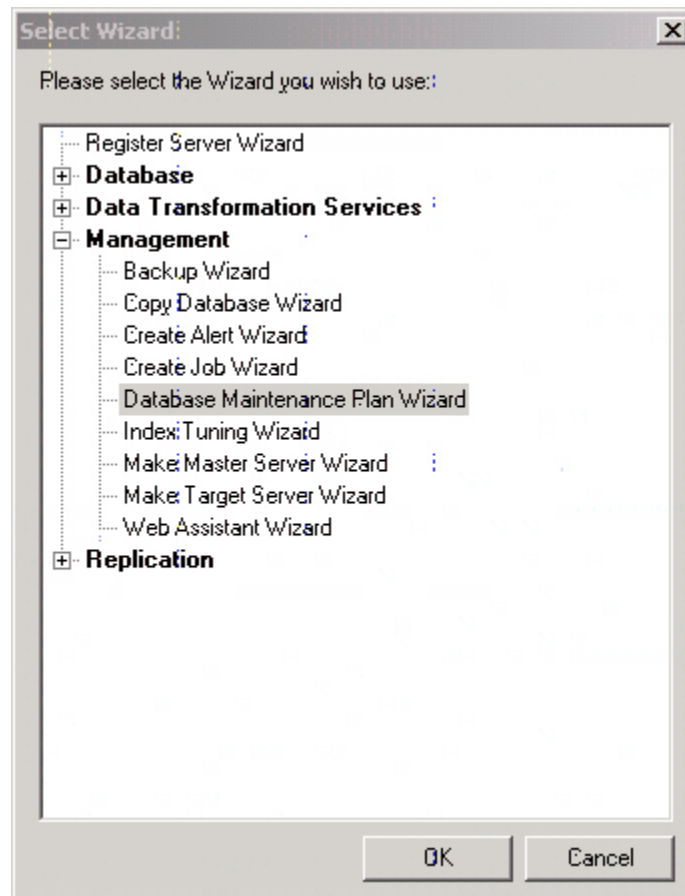
Log shipping

In SQL Server 2000 Enterprise Edition, log shipping can be used to send transaction logs from one database to another on a scheduled basis. This will allow synchronization of the source and the destination database. The use of this is very efficient with standby servers. Log shipping may be configured in Enterprise Manager. Complete instructions for standby server and log shipping are covered in SQL Server Books Online.

This example will use SQL Server 2000 Enterprise Edition to log ship transactions between two servers.

The initial procedure is the same as configuring a standby server. There should be a primary and standby server. The standby server should have the backups restored from the primary server and should be in read-only mode.

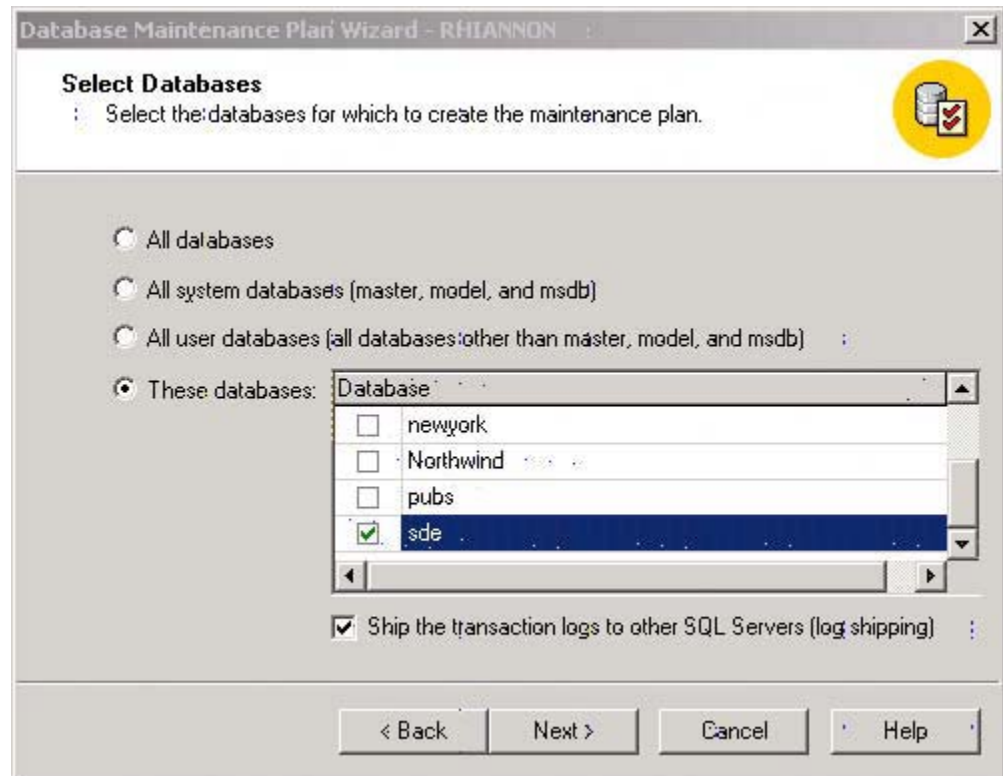
From the primary server, select Tools, Wizards, Management, Data Maintenance Plan Wizard.



Click OK to start the wizard.



Click Next>.



Select the database that the transaction logs are to be shipped from. Click the Ship the transaction logs to other SQL Servers check box. Click Next>.

Database Maintenance Plan Wizard - RHIANNON

Update Data Optimization Information

As data and index pages fill, updating requires more time. Reorganize your data and index pages to improve performance.

☒ Reorganize data and index pages

☐ Reorganize pages with the original amount of free space

☒ Change free space per page percentage to: %

☐ Update statistics used by query optimizer. Sample: % of the database

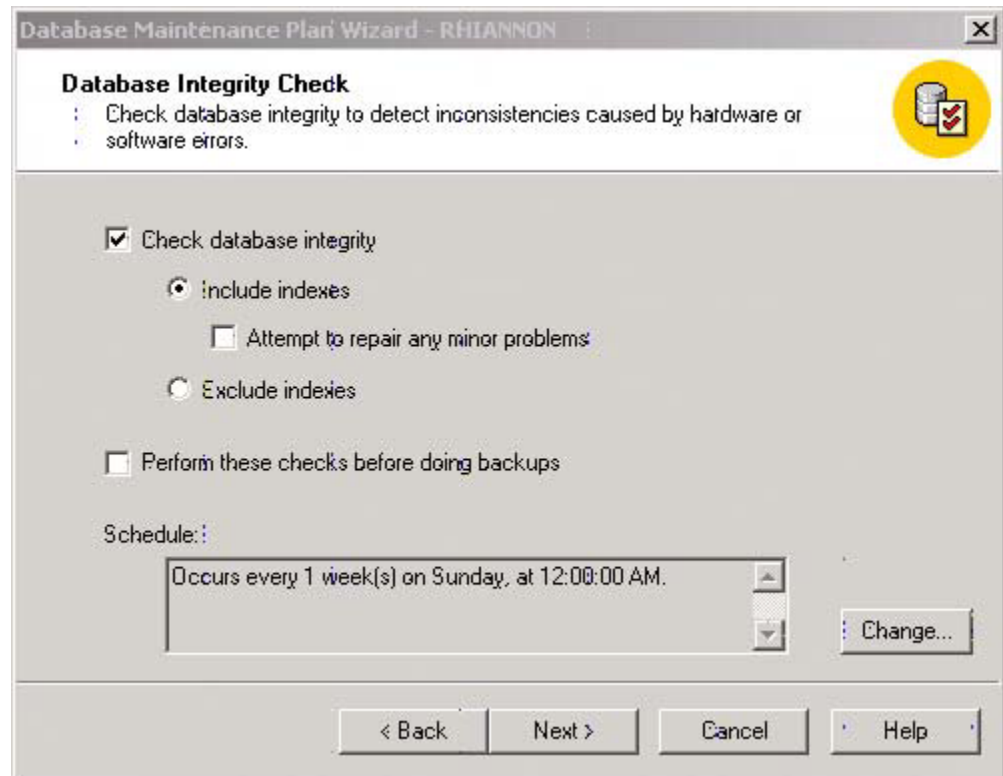
☒ Remove unused space from database files

When it grows beyond: MB

Amount of free space to remain after shrink: % of the data space

Schedule:

As part of the update Wizard, tasks other than log shipping are performed. All defaults for these tasks will be accepted. Appropriate entries should be made for your installation. Click Next>.



Default selections are shown. Selecting Attempt to repair any minor problems will put the database in single-user mode so that availability will be restricted. Click Next>.

The screenshot shows the 'Specify the Database Backup Plan' step of the 'Database Maintenance Plan Wizard - RHIANNON'. The window title is 'Database Maintenance Plan Wizard - RHIANNON'. The main heading is 'Specify the Database Backup Plan:' with a sub-instruction: 'Specify the database backup plan to prevent data loss due to system failure.' There is a yellow circular icon with a database cylinder and a checkmark in the top right corner. The options are as follows:

- ☒ Back up the database as part of the maintenance plan
- ☒ Verify the integrity of the backup when complete

Location to store the backup file:

- ☐ Tape: [Empty text box]
- ☒ Disk

Schedule:

Occurs every 1 week(s) on Sunday, at 2:00:00 AM.

[Change...]

Navigation buttons at the bottom: < Back, Next >, Cancel, Help.

Select the database backup options and click Next>.

The screenshot shows the 'Specify Backup Disk Directory' step of the 'Database Maintenance Plan Wizard - RHIANNON'. The window title is 'Database Maintenance Plan Wizard - RHIANNON'. The main heading is 'Specify Backup Disk Directory' with a sub-instruction: 'Specify the directory in which to store the backup file.' There is a yellow circular icon with a database cylinder and a checkmark in the top right corner. The options are as follows:

Directory in which to store the backup file:

- ☒ Use the default backup directory
- ☐ Use this directory: [C:\Program Files\Microsoft SQL Se] ...

☐ Create a subdirectory for each database

☐ Remove files older than: 4 [Week(s)]

Backup file extension: BAK

Navigation buttons at the bottom: < Back, Next >, Cancel, Help.

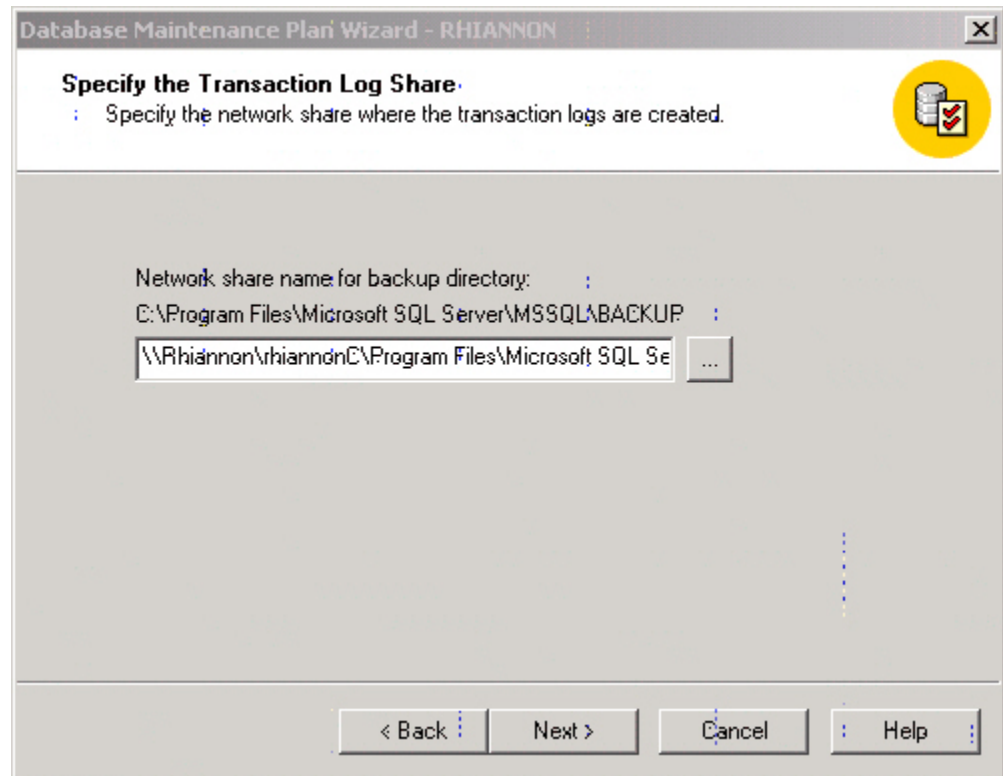
Specify the location to store the backup file. Click Next>.

The screenshot shows the 'Database Maintenance Plan Wizard - RHIANNON' window. The title bar includes the text 'Database Maintenance Plan Wizard - RHIANNON' and a close button. The main window has a title bar with 'Specify Transaction Log Backup Disk Directory' and a yellow circular icon with a database cylinder and a checkmark. Below the title bar, the instruction 'Specify the directory in which to store the transaction log backup file.' is displayed. The main area contains the following options:

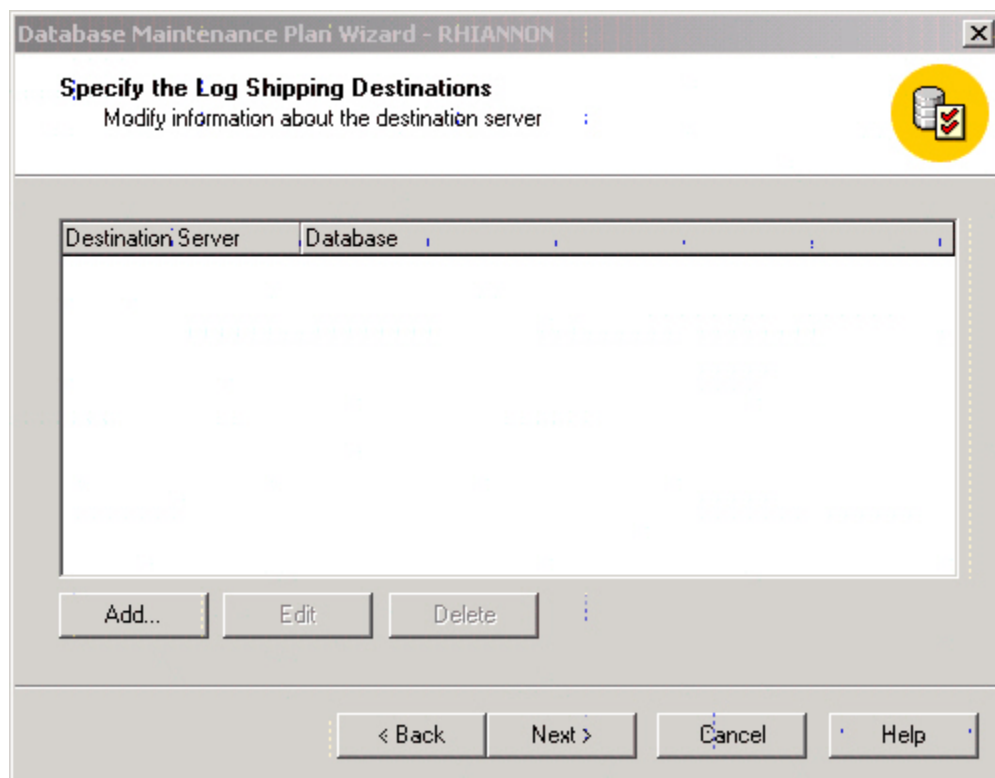
- 'Directory in which to store the backup file:' with two radio buttons:
 - ☒ Use the default backup directory
 - ☐ Use this directory: [C:\Program Files\Microsoft SQL Se] ...
- ☐ Create a subdirectory for each database
- ☐ Remove files older than: [4] [Week(s)]
- Backup file extension: [TRN]

At the bottom, there are four buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

Specify the location to store the transaction backup file. Click Next>.



In order to run log shipping, a share on the primary server drive that stores the transaction log must be created. Enter the location of this drive and click Next>.



Click Add to specify the log shipping destination server.

Add Destination Database

General

Server Name: RHIANNON\LOGSHP

Transaction Log Destination Directory:

Directory: C:\Program Files\Microsoft SQL Serve

Destination Database

☐ Create and initialize new database

Database name: sde

Use these file directories:

For data: C:\Program Files\Microsoft SQL Serve

For log: C:\Program Files\Microsoft SQL Serve

☒ Use existing database (No initialization)

Database Name: sde

Database Load State

☐ No recovery mode ☒ Standby mode

☐ Terminate users in database (Recommended)

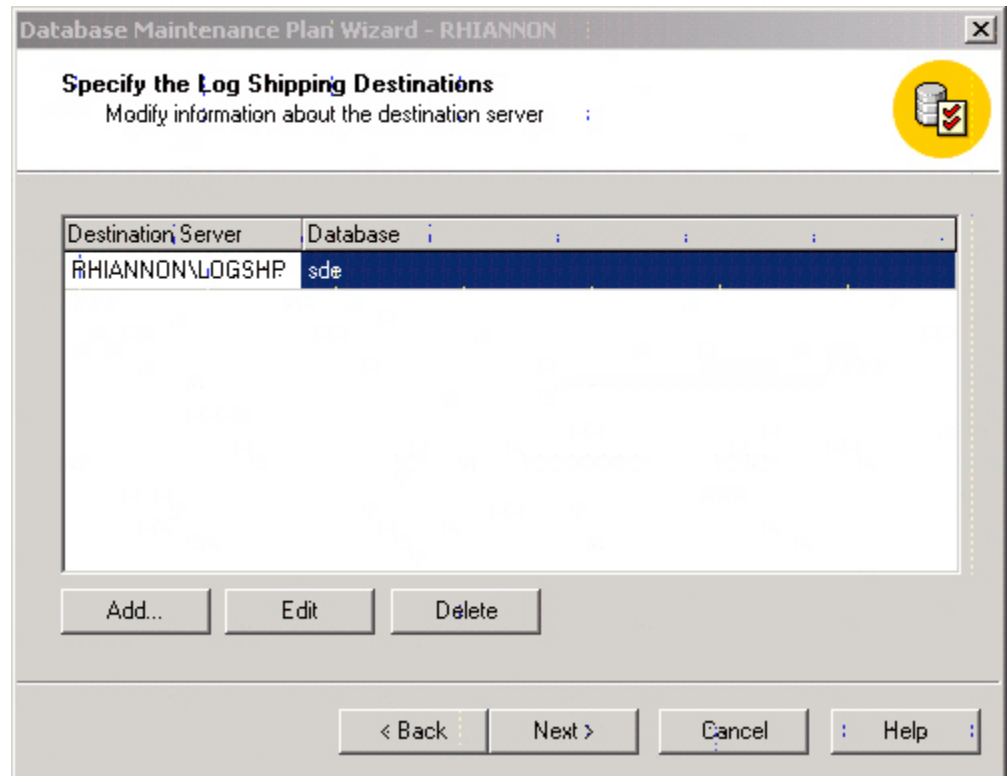
☐ Allow database to assume primary role

Transaction Log Backup Directory:

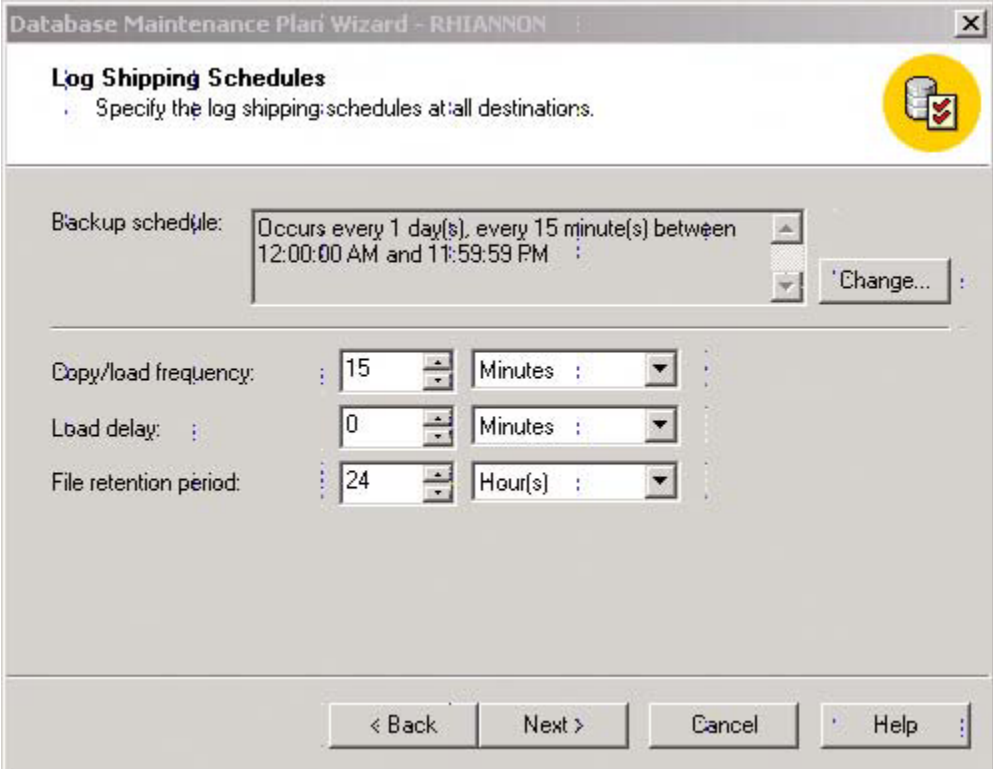
Directory:

OK Cancel Help

Enter the server name for the destination of the transaction logs. Here an existing database was used to receive the transaction logs. Click OK.



Confirm that the destination server and database are correct and click Next>.



The screenshot shows the 'Log Shipping Schedules' step of the 'Database Maintenance Plan Wizard - RHIANNON'. The window has a title bar with the name and a close button. Below the title bar is a header area with the title 'Log Shipping Schedules' and a sub-instruction 'Specify the log shipping schedules at all destinations.' To the right of the instruction is a yellow circular icon containing a database cylinder and a checkmark. The main area of the wizard contains several configuration options: 'Backup schedule:' with a text box showing 'Occurs every 1 day(s), every 15 minute(s) between 12:00:00 AM and 11:59:59 PM' and a 'Change...' button; 'Copy/load frequency:' with a spinner set to '15' and a dropdown set to 'Minutes'; 'Load delay:' with a spinner set to '0' and a dropdown set to 'Minutes'; and 'File retention period:' with a spinner set to '24' and a dropdown set to 'Hour(s)'. At the bottom of the window are four buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

Database Maintenance Plan Wizard - RHIANNON

Log Shipping Schedules

Specify the log shipping schedules at all destinations.

Backup schedule: Occurs every 1 day(s), every 15 minute(s) between 12:00:00 AM and 11:59:59 PM Change...

Copy/load frequency: 15 Minutes

Load delay: 0 Minutes

File retention period: 24 Hour(s)

< Back Next > Cancel Help

Enter the backup schedule for the log shipping and click Next>.

Database Maintenance Plan Wizard - RHIANNON

Log Shipping Thresholds

Specify thresholds that, if exceeded, will generate alerts.

Backup alert threshold: 45 Minutes

Out of sync alert threshold: 45 Minutes

< Back Next > Cancel Help

Specify the log shipping thresholds and click Next>.

The screenshot shows a Windows-style dialog box titled "Database Maintenance Plan Wizard - RHIANNON". The window has a standard title bar with a close button (X) in the top right corner. Below the title bar, the main content area has a light gray background. At the top of this area, the text "Specify the Log Shipping Monitor Server Information" is displayed in bold. Below this, a subtitle reads "Specify the central server from which log shipping will be monitored." To the right of this text is a yellow circular icon containing a database cylinder and a checkmark. The main configuration area contains a "SQL Server:" label followed by a dropdown menu currently showing "RHIANNON". Below this are two radio button options: "Use Windows Authentication" (which is selected) and "Use SQL Server Authentication". Under the "Use SQL Server Authentication" option, there are two text input fields: "Login name:" with the text "log_shipping_monitor_probe" and "Password:" with an empty field. At the bottom of the window, there is a gray bar containing four buttons: "< Back", "Next >", "Cancel", and "Help".

Specify the server to monitor the log shipping. Click Next>.

Database Maintenance Plan Wizard - RHIANNON

Reports to Generate

: Specify the directory in which to store the reports generated by the maintenance plan.

☐ Write report to a text file in directory: C:\Program Files\Microsoft SQL Se...

☐ Delete text report files older than: 4 Week(s)

☐ Send e-mail report to operator: ...

New Operator...

< Back Next > Cancel Help

Specify any reports to generate. Click Next>.

The screenshot shows the 'Maintenance Plan History' step of the 'Database Maintenance Plan Wizard - RHIANNON'. The window title is 'Database Maintenance Plan Wizard - RHIANNON'. The main heading is 'Maintenance Plan History' with a sub-instruction: 'Specify how you want to store the maintenance plan records.' There is a yellow circular icon with a database cylinder and a checkmark in the top right corner.

The 'Local server' section is expanded, showing the following options:

- ☒ Write history to the msdb.dbo.sysdbmaintplan_history table on this server
- ☒ Limit rows in the table to: rows for this plan

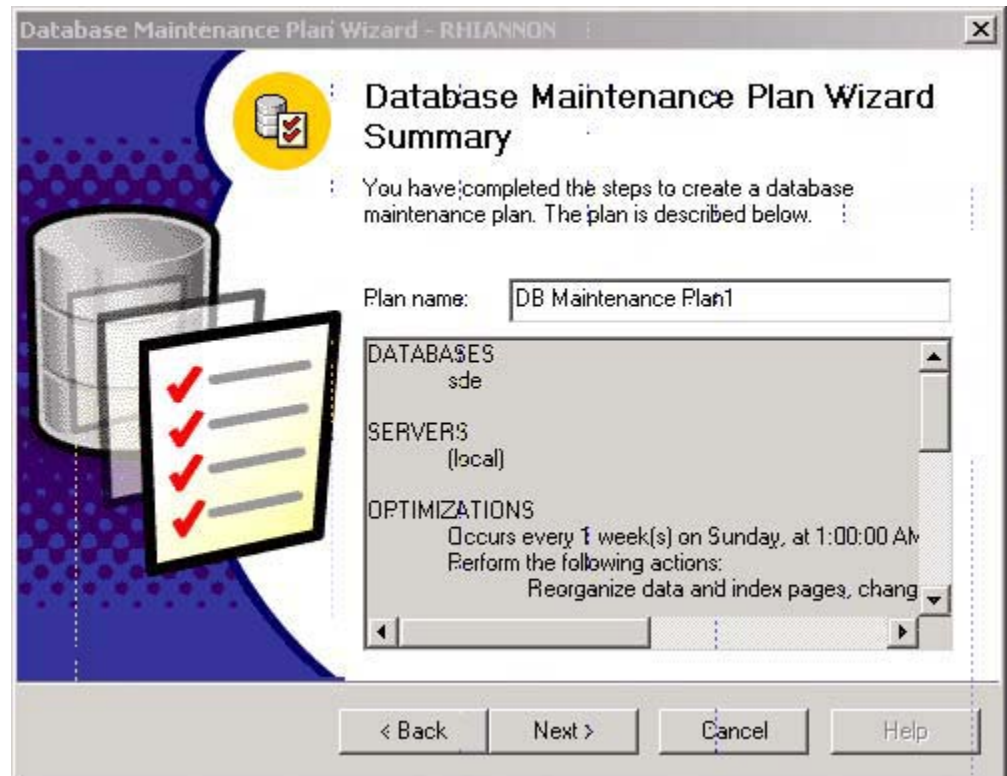
The 'Remote server' section is collapsed. Below it, a note states: 'History is added to the msdb.dbo.sysdbmaintplan_history table on the remote server. : Windows Authentication is used to log on to the remote server.'

Below the note, the 'Remote server' section is expanded, showing the following options:

- ☐ Write history to the server: ...
- ☒ Limit rows in the table to: rows for this plan

At the bottom of the wizard, there are four buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

Specify how to store the maintenance plan records. Click Next>.



Name the database maintenance plan. Here the default DB Maintenance Plan1 was accepted. Confirm the options selected and click Next>.



Click Finish to create a database maintenance plan.

It will be necessary to start the SQL Server Agent. In Enterprise Manager, open Management and right-click SQL Server Agent to start. The log shipping will be performed at the times specified in the log shipping schedules. The steps to restore the standby server are the same as above. If the production server should become unavailable, apply all remaining transaction logs to the standby server. From Query Analyzer recover the standby server database without restoring. Execute the following statement with the name of the database to be recovered.

```
-- Restore database using WITH RECOVERY.
RESTORE DATABASE sde
    WITH RECOVERY
```

The standby server will now be in the same state as the production server from the last transaction log. All changes after the last backup of the transaction log will be lost.

CHAPTER 8

Backup and Recovery

Introduction

MS SQL Server allows you to backup a database from one server and restore it to another. You can also detach and reattach databases between servers. This functionality allows you an easy way of distributing your data to other servers without having to worry about reloading it and rebuilding relationships or networks.

If you restore a database backup from one server to another, the code pages and sort orders of those two servers must be identical. For example, you cannot restore a database built with Binary Sort Order (BSO) to a server configured for default (dictionary) sort order. You can ascertain your server's sort order by running the stored procedure 'sp_helpsort' in the master database.

In this example, this server is configured for default, or dictionary, sort order.

```
Server default collation
Latin1-General, case-insensitive, accent-sensitive, kanatype-insensitive,
width-insensitive for Unicode Data, SQL Server Sort Order 52 on Code Page
1252 for non-Unicode Data
```

A backup operation also backs up the database's system tables including sysusers. When you apply a backup to a different server, the users listed in the database may differ from those users listed in the new server's syslogins table in the master database. This will prevent you from logging into the restored database until you correct the login mapping between the master database's syslogins table and the restored database's sysusers table. To synchronize the two tables, run sp_change_users_login.

Moving data using backup and restore

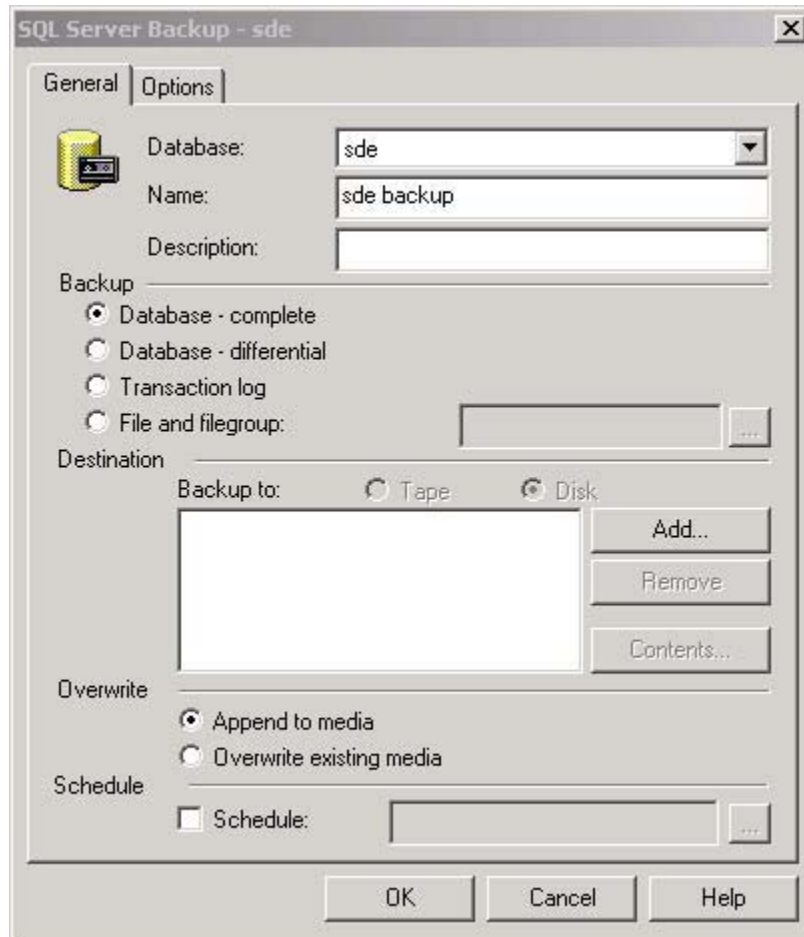
This section assumes you are going to backup an ArcSDE database and apply that backup to another server that does not have a preexisting ArcSDE installation.

1. Backup the SDE database and any other database you wish to apply to another server.
2. Copy the backup file(s) to the new server.
3. Restore the backup file(s) to the new server.
4. Run sp_change_users_login to synchronize the sysuser's table in each newly restored database with the master database's syslogins table.


1. Backup the SDE database

Use either t-sql's BACKUP command or the Enterprise Manager's backup utility to backup the sde and any other spatial database you have. In this example, we use the Enterprise Manager to backup the sde database.

Expand Enterprise Manager's Server Group, then your server host name. Right-click your database, select all tasks backup database.



The backup database form will appear. You'll need to add a destination with the 'add' button if none appears. Click Add and choose a directory to backup to.



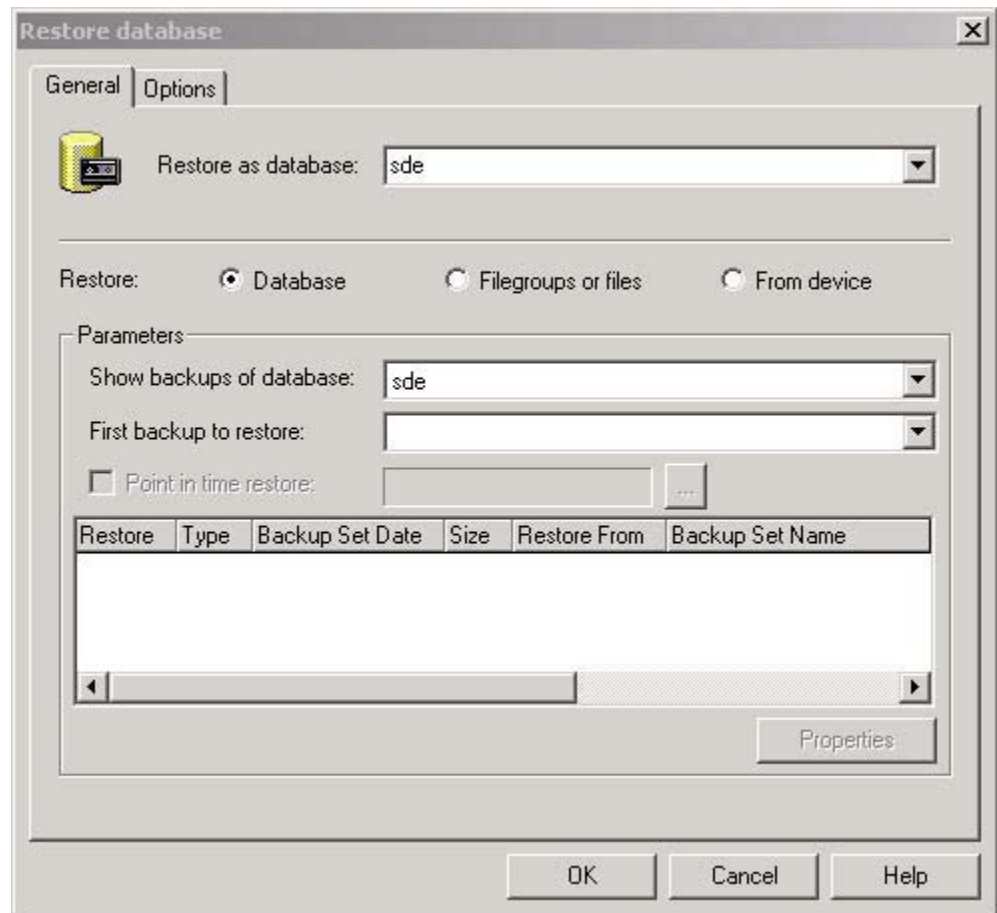
Click OK when you are done. See the SQL Server Books Online for exact information on backup types and options.

2. Copy the backup file(s) to the new server

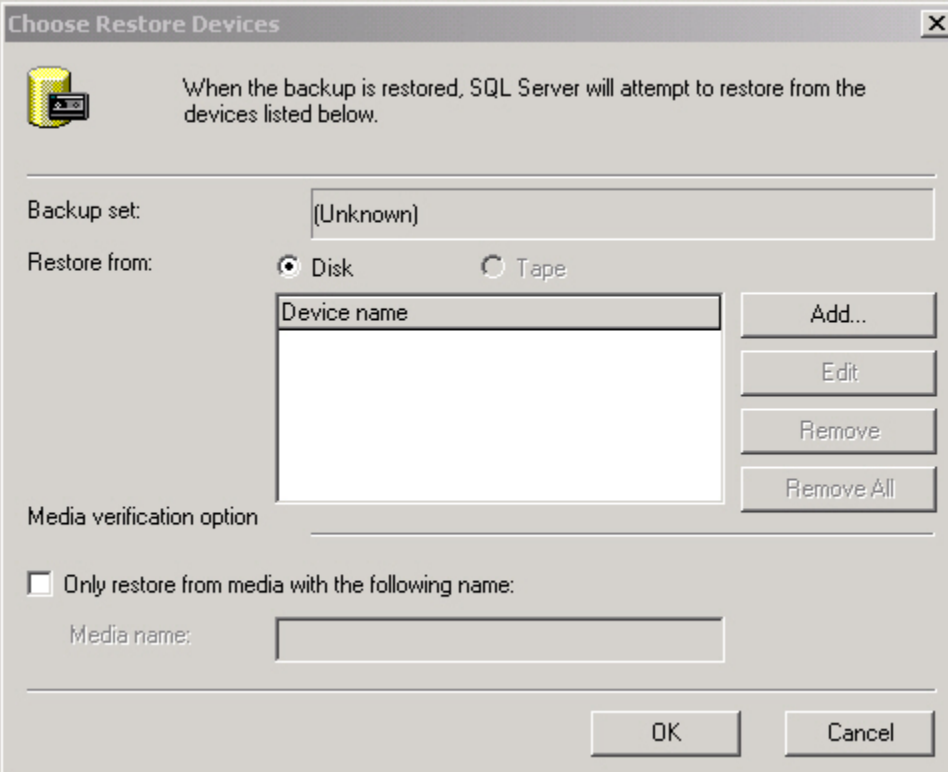
In the example above, the backup file is located on D:\rhannon_backup. Copy the backup file to the new server.

3. Restore the backup file(s) to the new server

Expand the new server in the Enterprise Manager's server group. Right-click 'databases' and select all tasks restore database.



The Restore database dialog box will appear. Click the From device radio button and then the "Select Devices" button.



The dialog box is titled "Choose Restore Devices" and has a close button (X) in the top right corner. It contains a yellow tape icon and the text: "When the backup is restored, SQL Server will attempt to restore from the devices listed below."

Backup set: (Unknown)

Restore from: ☒ Disk ☐ Tape

Device name

Media verification option

☐ Only restore from media with the following name:

Media name:

Buttons: Add..., Edit, Remove, Remove All, OK, Cancel

If you want to restore the database with a different name, change its name in the Restore as database combo box. NOTE: If you do this, you may cause inconsistencies in the sde or gdb metadata tables. Navigate to the backup file and select it. Click OK all the way back to this form and select the Options tab.

Restore database

General Options

☐ Eject tapes (if any) after restoring each backup

☐ Prompt before restoring each backup

☐ Force restore over existing database

Restore database files as: Read from media

Logical file name	Move to physical file name
sde_dat	f:\ss2k_files\MSSQL\data\sdedata.mdf
sde_log	f:\ss2k_files\MSSQL\data\sdelog.ldf

Recovery completion state

☒ Leave database operational. No additional transaction logs can be restored.

☐ Leave database nonoperational but able to restore additional transaction logs.

☐ Leave database read-only and able to restore additional transaction logs.

Undo file: f:\ss2k_files\MSSQL\BACKUP\UNDO_sde.DAT

OK Cancel Help

In the middle of this form, you'll see a grid control with paths in it. You'll want to verify that the paths are correct for your restore file. You can restore these files to any location you like, as long as it exists on disk. If you check Force restore over existing database, you'll restore this database over an existing one. If you don't and you have an existing database with the same name, the restore will fail. Click OK.

Once the restore finishes, open the database in the Enterprise Manager and verify your tables are there. Now you must synchronize the database users with their logins located in master.dbo.sysxlogins.

4. Run sp_change_users_login

When you restore a database from a backup created on another machine, there is a chance that a database user's identification number (either an SID or a GUID) does not match the same entry in the master.dbo.sysxlogins table. There is also a chance that the database user does not exist in the master.dbo.sysxlogins table. Since a restore also restores a database's sysusers table, it is possible for a record to exist here but not be present in master.dbo.sysxlogins. Hence, there are two scenarios you need to consider:

- € The user exists as a login on server but has a different identification (SID or GUID).

✎ The user does not exist (added as a login) on the server.

This does not apply to Windows Authenticated logins.

Scenario 1: The user exists on both servers but has a different identifier

In this example, both servers have a login called 'sde.' After you restore the sde database onto the new server, the sde user's SID differs in the sde database from the sde login's SID in the master.dbo.sysxlogins table.

```
Use sde
go
Select SID from sysusers where name = 'sde'
Use master
go
select SID from sysxlogins where name = 'sde'
```

```
SID
-----
0x76695419BFAED41184FD00C04F8D0451

(1 row(s) affected)
```

```
SID
-----
0xEDDFCA8E56B0D411850000C04F8D0451

(1 row(s) affected)
```

To synchronize the two accounts, run `sp_change_users_login`. For exact usage, see the SQL Server Books Online. In this example, we use the "update_one" option to update just the sde user. You must run `sp_change_users_login` within the database you restored (or want to change the login in).

```
Use sde
go
sp_change_users_login 'update_one','sde','sde'
```

Rerun the SID query above to verify your results:

```
0xEDDFCA8E56B0D411850000C04F8D0451
0xEDDFCA8E56B0D411850000C04F8D0451
```

Scenario 2: The user does not exist on the destination (restore) server

This scenario differs from the first one in that a user exists in the database but not as a login in the master.dbo.sysxlogins table. This can happen when you create a login on a server, backup the databases on that server, then restore the databases to another server that does not have a record for this login. In this example, the sde database is restored onto a machine that never had an 'sde' login.

```
Use sde
go
Select SID from sysusers where name = 'sde'
Use master
go
select SID from sysxlogins where name = 'sde'
```

```
SID
-----
0x76695419BFAED41184FD00C04F8D0451

(1 row(s) affected)
```

```
SID
```



```
-----
(0 row(s) affected)
```

To rectify this, you must add an 'sde' login to the server but not grant them access to the sde database. To add a login, use the `sp_addlogin` stored procedure or use the logins tool with the Enterprise Manager:

```
sp_addlogin 'sde','go','sde'
```

Now run `sp_change_users_login` to synchronize the SID values in `master.dbo.sysxlogins` and the restored database's `sysusers` table. You must run `sp_change_users_login` in the sde database since its `sysusers` table is out of sync with the `master.dbo.sysxlogins` table.

```
use sde
go
sp_change_users_login 'update_one','sde','sde'
```

Now verify that the SIDs are equal by rerunning our SID query.

```
Use sde
Go
Select SID from sysusers where name = 'sde'
Use master
Go
select SID from sysxlogins where name = 'sde'
SID
```

```
-----
0xF6DFCA8E56B0D411850000C04F8D0451
```

```
(1 row(s) affected)
```

```
SID
```

```
-----
0xF6DFCA8E56B0D411850000C04F8D0451
```

```
(1 row(s) affected)
```

Once you've synchronized the users, you are ready to either start the ArcSDE service or make direct connections to it.

Moving Data using sp_detach_db and sp_attach_db

You can also move databases using the SQL Server stored procedures `sp_detach_db` and `sp_attach_db`. A detached database can be one file or several files, depending upon how you created the database. The procedure is similar to the backup and restore example, but differs in that you don't have a backup file of your database. Be careful – once you attach your database, you don't have a copy of it, so if you delete it, you'll lose your data.

To move your database(s) with `sp_detach_db` and `sp_attach_db`:

1. Detach your database from one server.
2. Copy the datafiles and transaction logfiles to the target server and re-attach the database there with `sp_attach_db`.
3. Run `sp_change_users_login` to synchronize the `sysuser`'s table in each newly restored database with the master database's `sysxlogins` table (SQL Server logins only).

1. Detach the database

Read the SQL Server Books Online for a complete discussion of `sp_detach_db`. When you detach the database, you unhook it from SQL Server without having to shut SQL Server down. For example:

```
use master
sp_detach_db 'sde'
go
```

2. Copy the Database to Another Server and Run sp_attach_db

The previously detached database will reside in its previous data folder. Copy it and its log file to the target server. Then execute the procedure `sp_attach_db` to reattach the database to the server:

```
sp_attach_db 'sde',
@filename1='d:\SQL2K_DATAFILES\MSSQL\Data\sdedata.mdf',
@filename2='d:\SQL2K_DATAFILES\MSSQL\Data\sdelog.ldf'
```

You must supply a pathname for the database files and logfiles.

2. Run sp_change_users_login to synchronize the sysuser's table

Follow the same procedure outlined above in step 4.

CHAPTER 9

Configuring snapshot replication

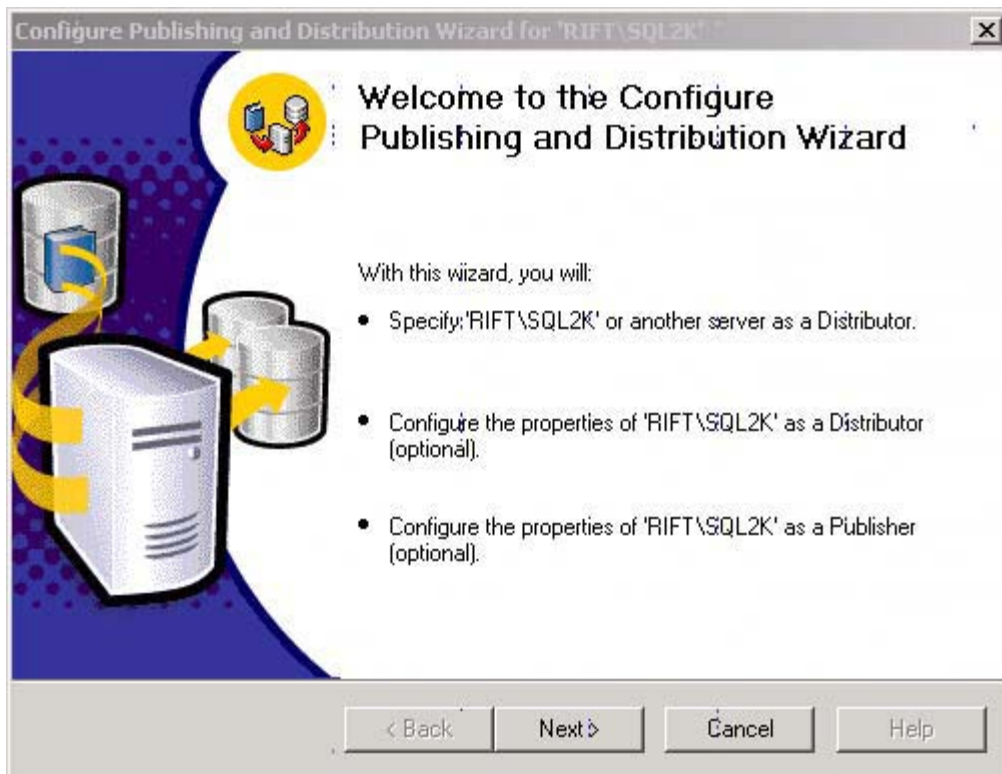
Replication is used to create multiple copies of the same data that can be distributed to multiple other databases throughout an organization.

Implementing snapshot replication

Snapshot replication is used to distribute an entire copy of the information from the distributing database to the subscribing database. In SQL Server 7 or SQL Server 2000, it is the easiest type of replication to setup and maintain. It is most appropriate for read-only data, and there is a high degree of latency in this process. Any changes that are made to the data at the subscriber will be overwritten once the next snapshot image is downloaded from the distributor. Management of the distribution of the snapshot can be performed from the distribution server, a push subscription, or from the subscription server, a pull subscription. The use of a push or pull subscription is dependent on the particular circumstances of each organization. A push subscription can be centrally managed but will require a greater storage premium on the distribution server.

Implementation of snapshot replication, in this example replication, will be between two instances of SQL Server 2000. The procedure will be similar for SQL Server 7.


From Enterprise Manager, for the SQL Server instance that you wish to replicate from, choose **TOOLS; Replication; and Configure Publishing, Subscribers, and Distribution**. This will begin the process of configuring the distribution server.



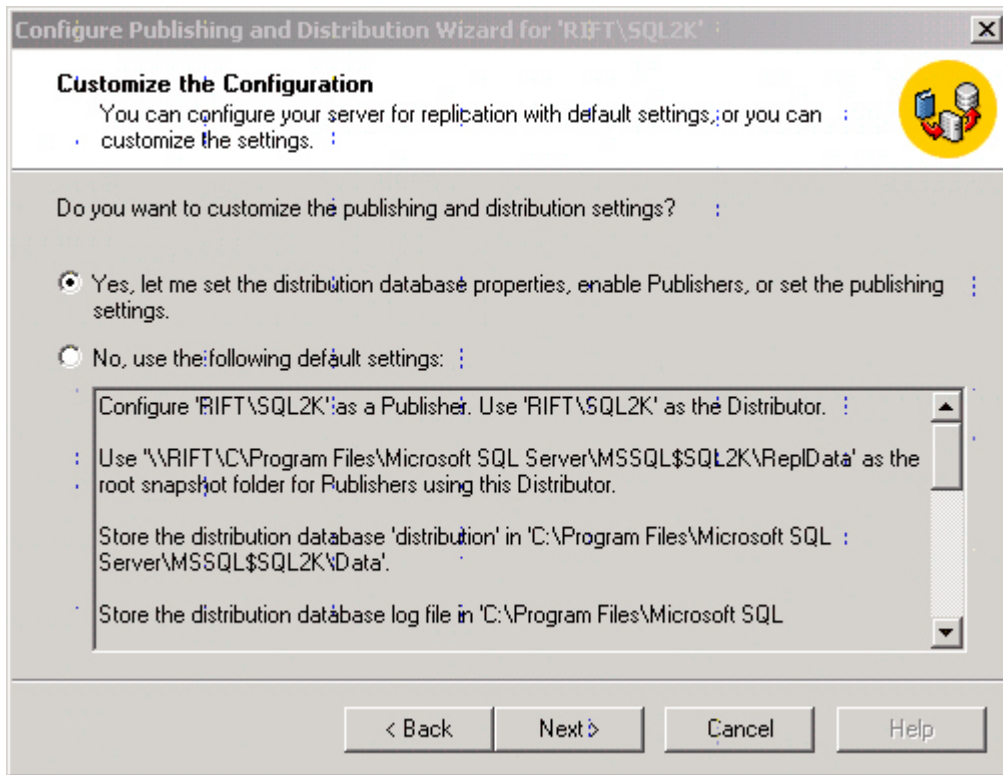
Click Next>.



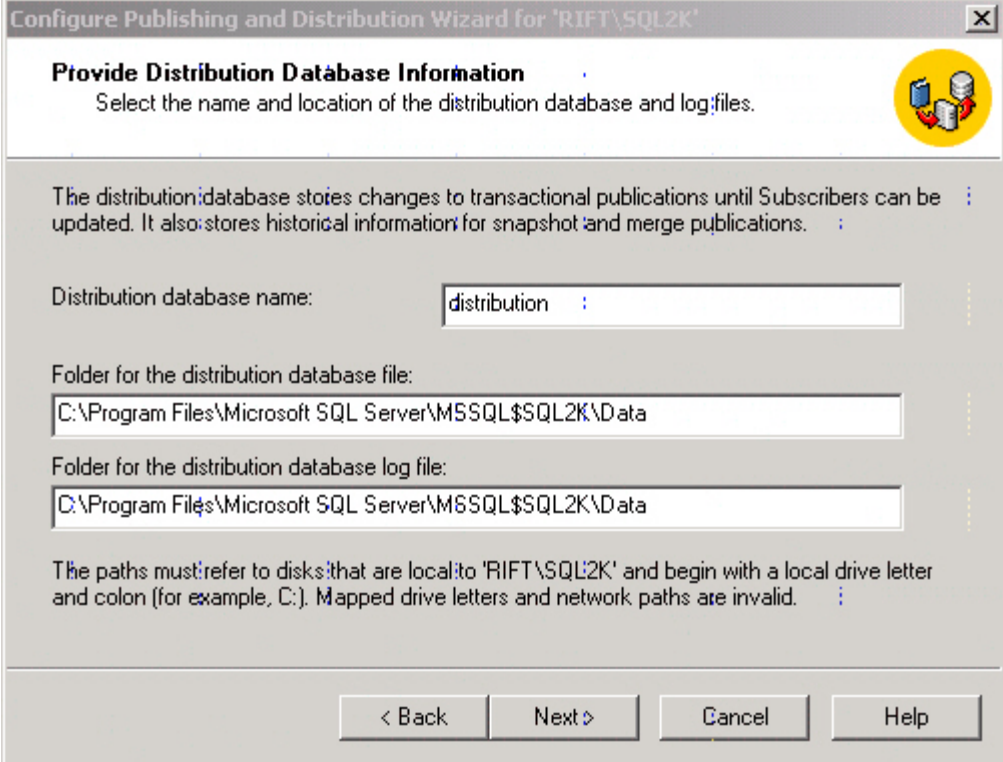
Choose the server to act as the distributor, then click Next>.



Specify the path to the replication snapshot folder, then choose Next>.



You may accept the default configuration from SQL Server, but for this example we will choose to customize the publishing and distribution settings. Click Next>.



The screenshot shows a Windows-style dialog box titled "Configure Publishing and Distribution Wizard for 'RIFT\SQL2K'". The window has a standard title bar with a close button (X) in the top right corner. The main content area is titled "Provide Distribution Database Information" and includes a sub-instruction: "Select the name and location of the distribution database and log files." To the right of this text is a yellow circular icon containing a database cylinder and a document with a red arrow. Below the instruction, there is a paragraph explaining the distribution database's role: "The distribution database stores changes to transactional publications until Subscribers can be updated. It also stores historical information for snapshot and merge publications." Following this, there are three text input fields. The first is labeled "Distribution database name:" and contains the text "distribution". The second is labeled "Folder for the distribution database file:" and contains the path "C:\Program Files\Microsoft SQL Server\MSSQL\$SQL2K\Data". The third is labeled "Folder for the distribution database log file:" and contains the same path "C:\Program Files\Microsoft SQL Server\MSSQL\$SQL2K\Data". Below these fields is another paragraph: "The paths must refer to disks that are local to 'RIFT\SQL2K' and begin with a local drive letter and colon (for example, C:). Mapped drive letters and network paths are invalid." At the bottom of the dialog, there are four buttons: "< Back", "Next >", "Cancel", and "Help".

Configure Publishing and Distribution Wizard for 'RIFT\SQL2K'

Provide Distribution Database Information
Select the name and location of the distribution database and log files.

The distribution database stores changes to transactional publications until Subscribers can be updated. It also stores historical information for snapshot and merge publications.

Distribution database name:

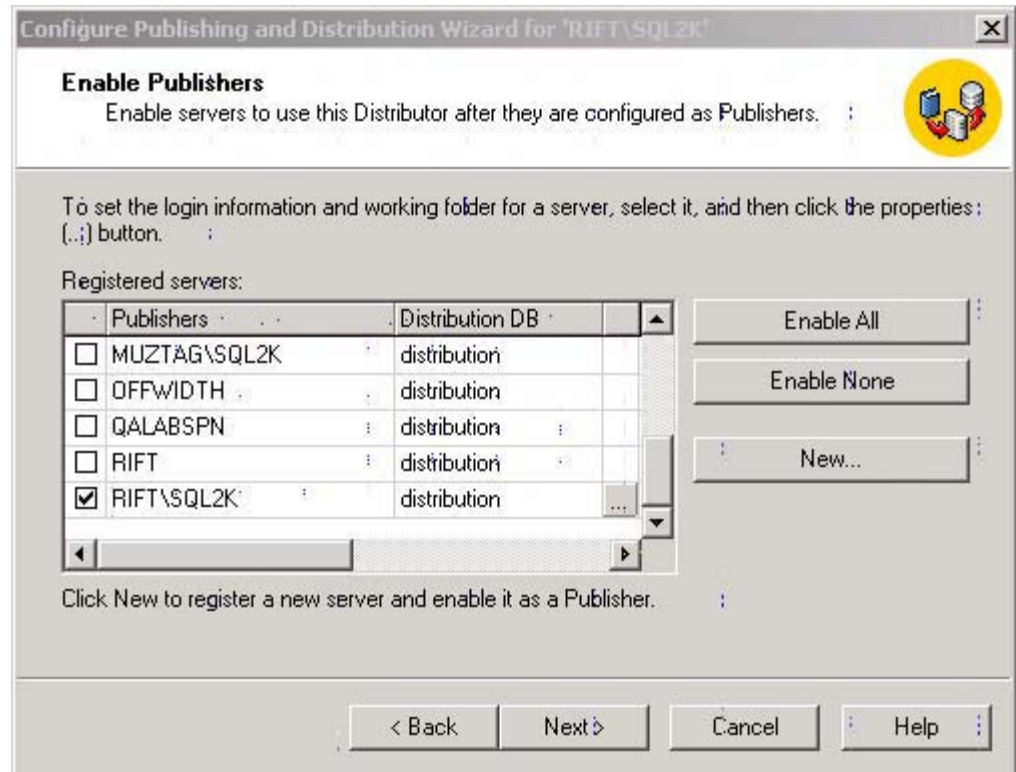
Folder for the distribution database file:

Folder for the distribution database log file:

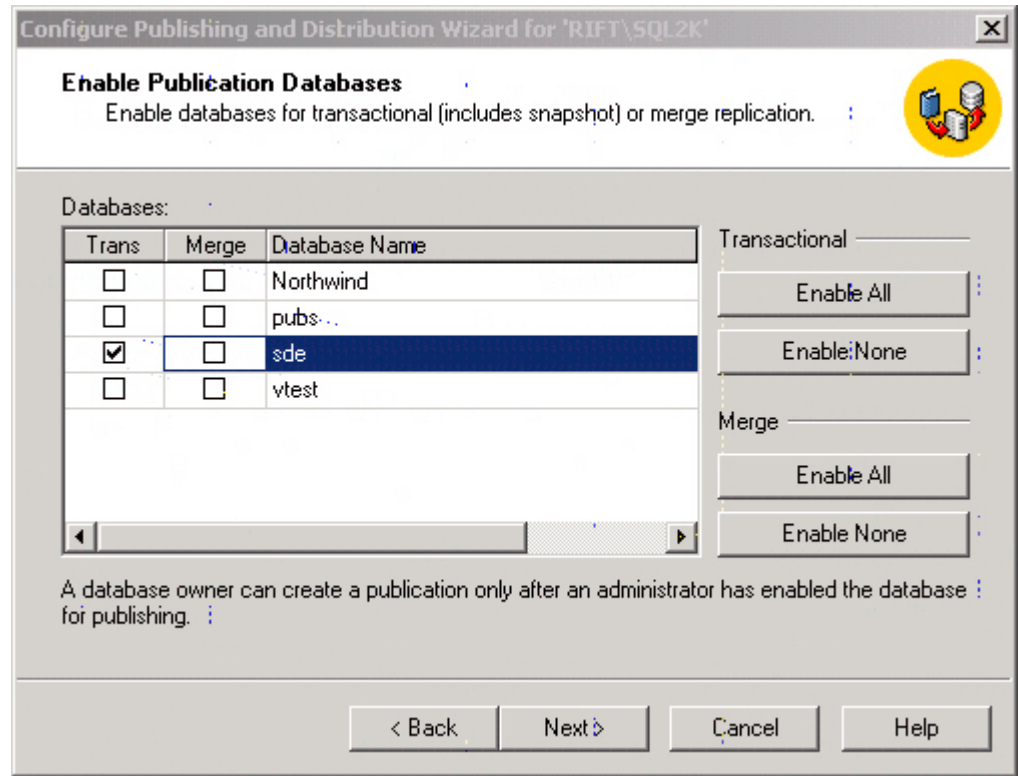
The paths must refer to disks that are local to 'RIFT\SQL2K' and begin with a local drive letter and colon (for example, C:). Mapped drive letters and network paths are invalid.

< Back Next > Cancel Help

Select the name and location of the distribution database and logfiles. Then choose Next>.



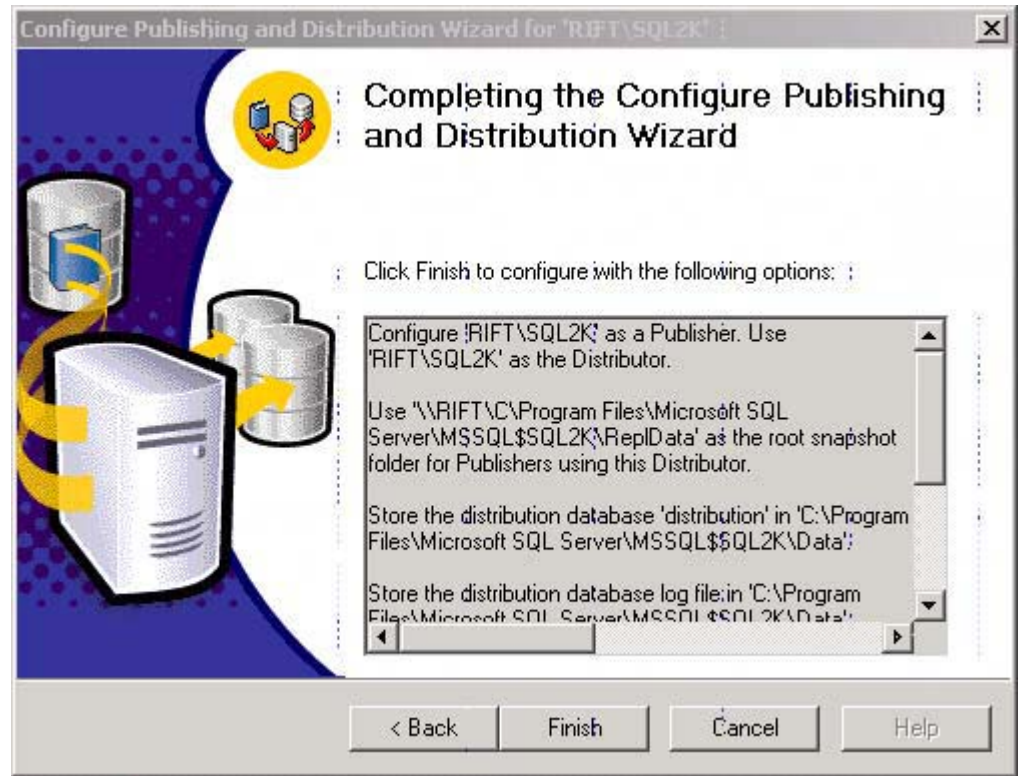
Enable the servers to use this distributor after they are configured as Publishers. Then choose Next>.



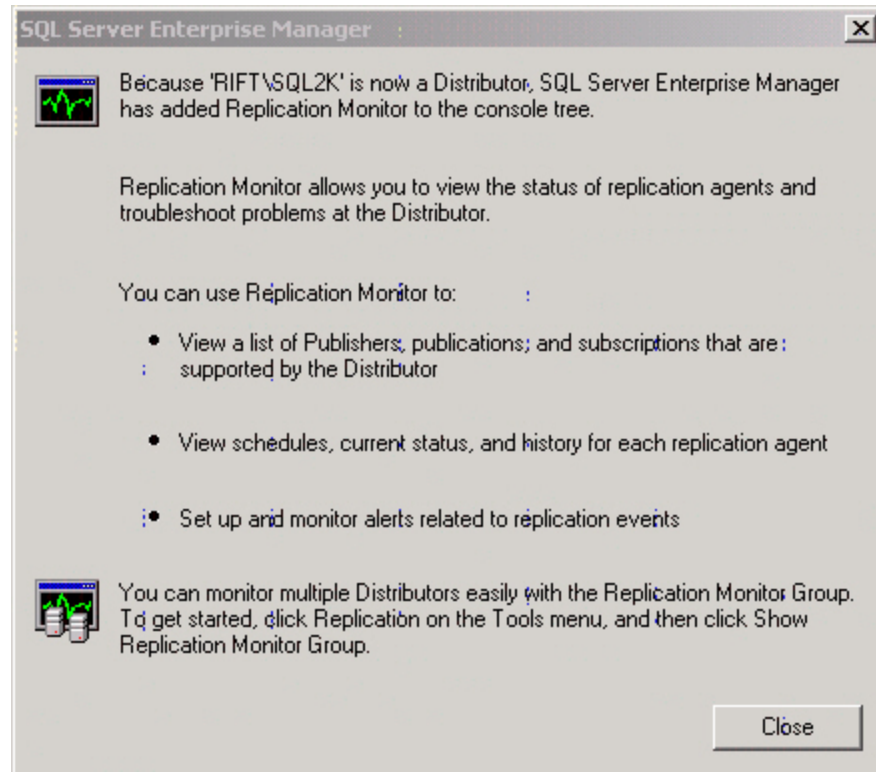
Enable the databases that are to be used for replication. Here choose Transactional, which includes snapshot. Choose Next>.



Enable the subscribing server to receive publications from the distributor. Then choose Next>.

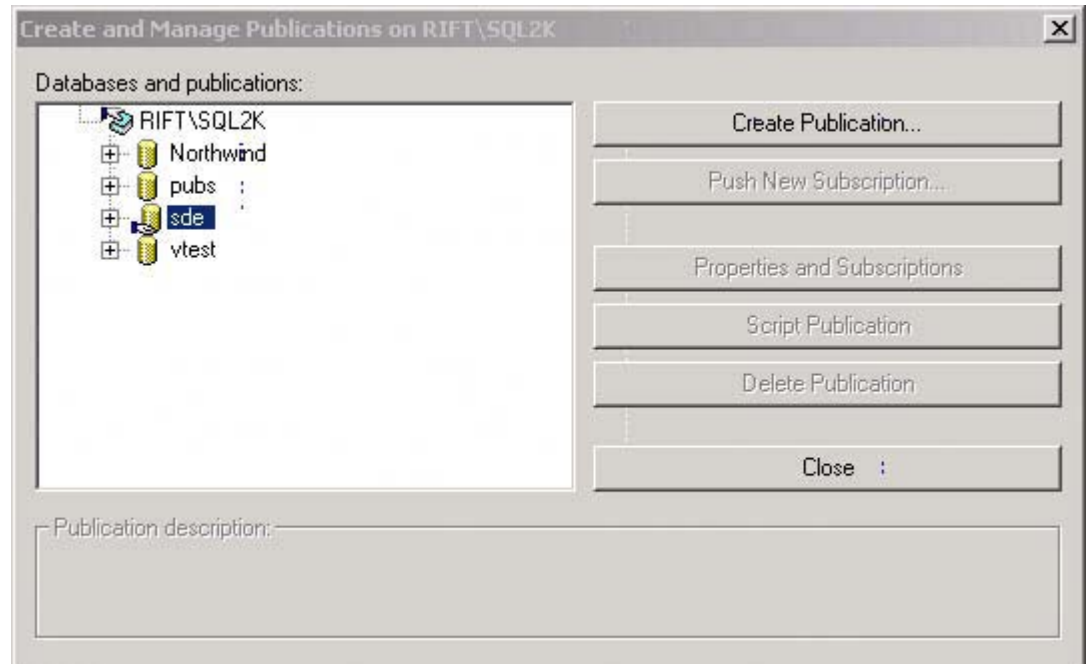


Choose Finish to complete the configuration of the distribution server.

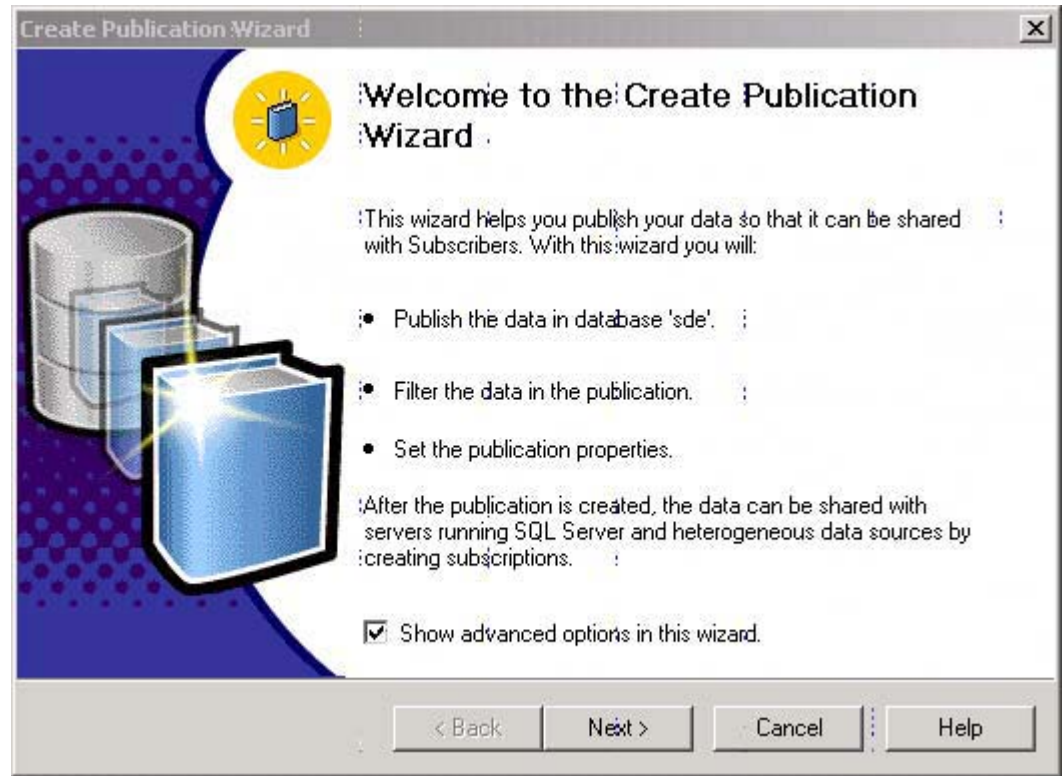


Choose Close to complete the distribution configuration process.

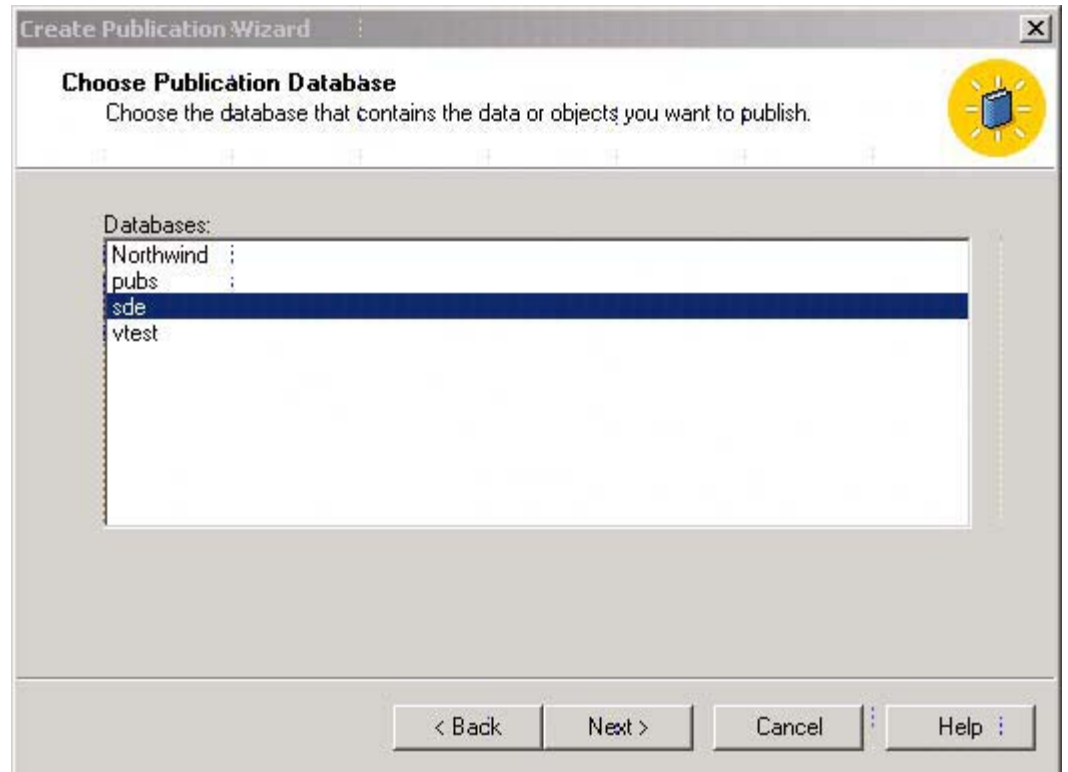
To create and manage a publication in Enterprise Manager, choose Replication, then choose Create and manage publications.



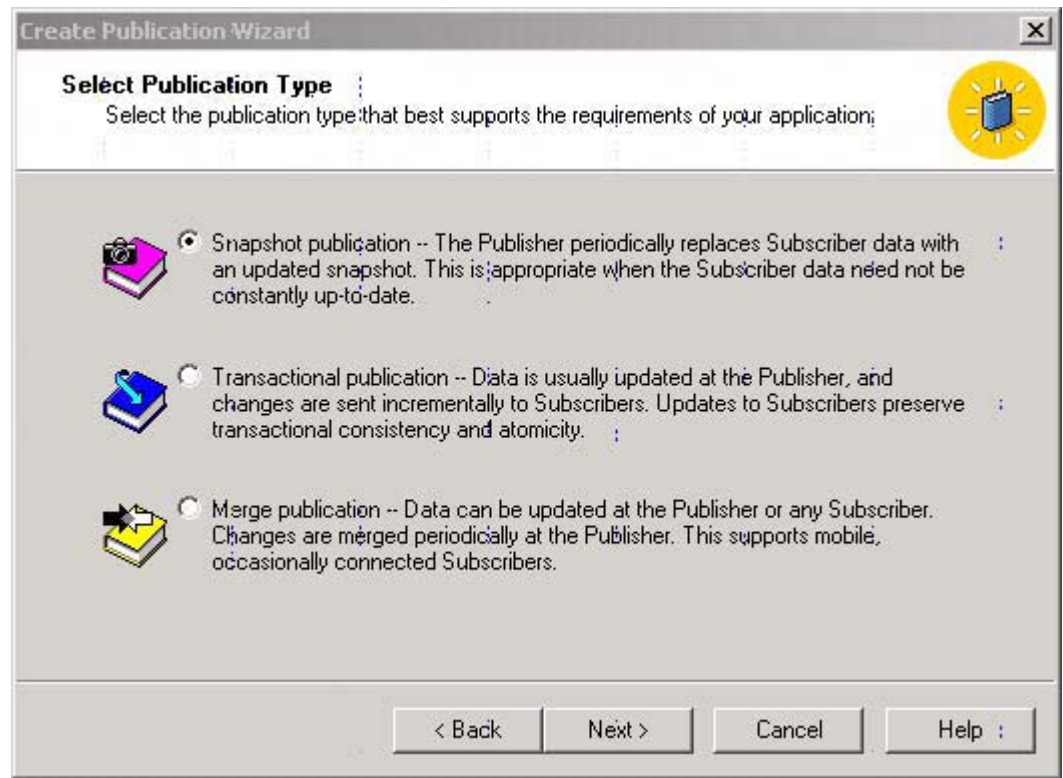
Highlight the source of the publication and click Create Publication.



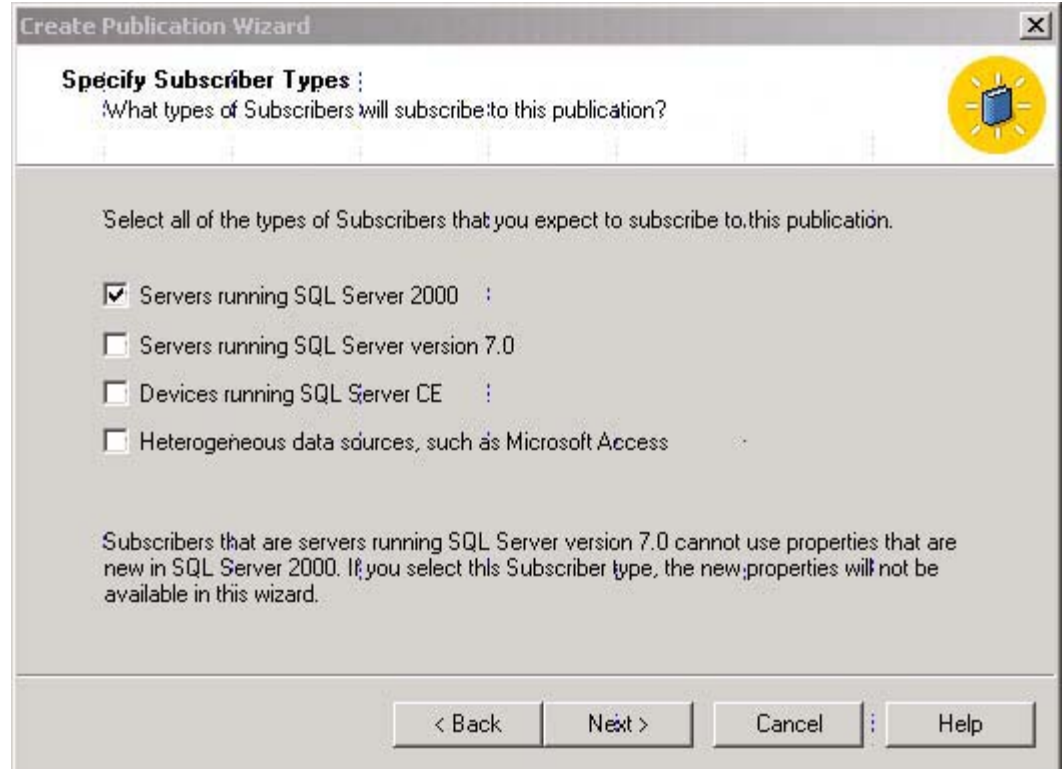
The Welcome to the Create Publication Wizard will appear. We will run this example with the Show advanced options in this wizard configured. Choose Next> to continue.



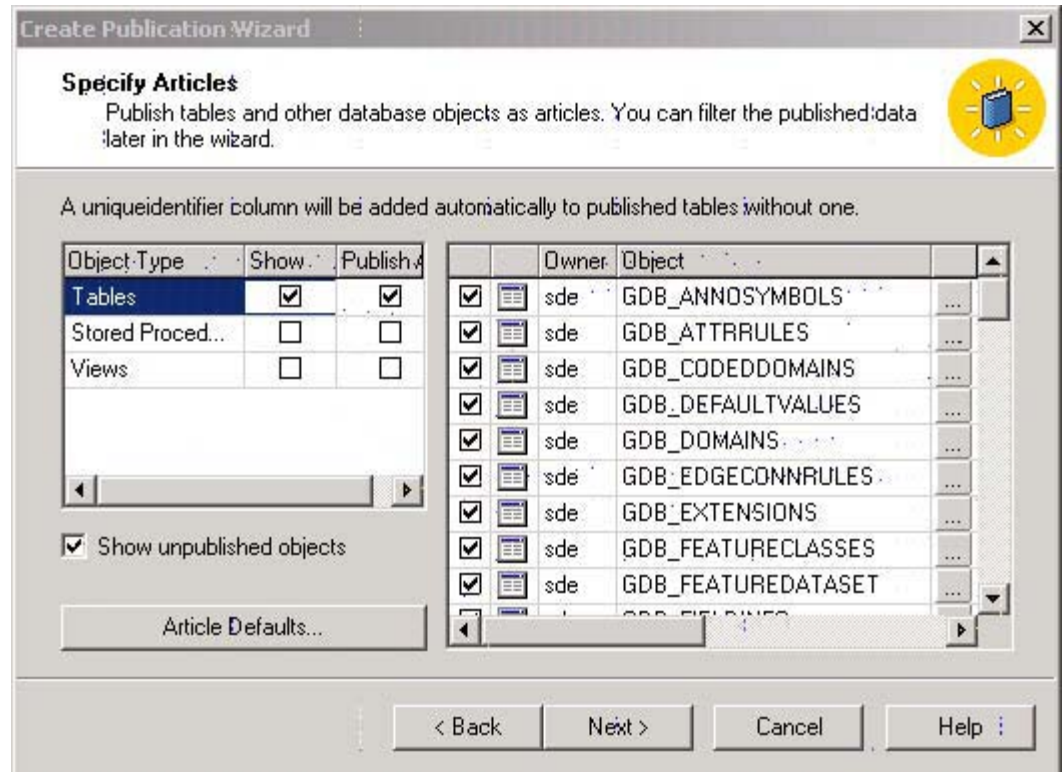
Highlight the database that contains the objects to be published, then click Next>.



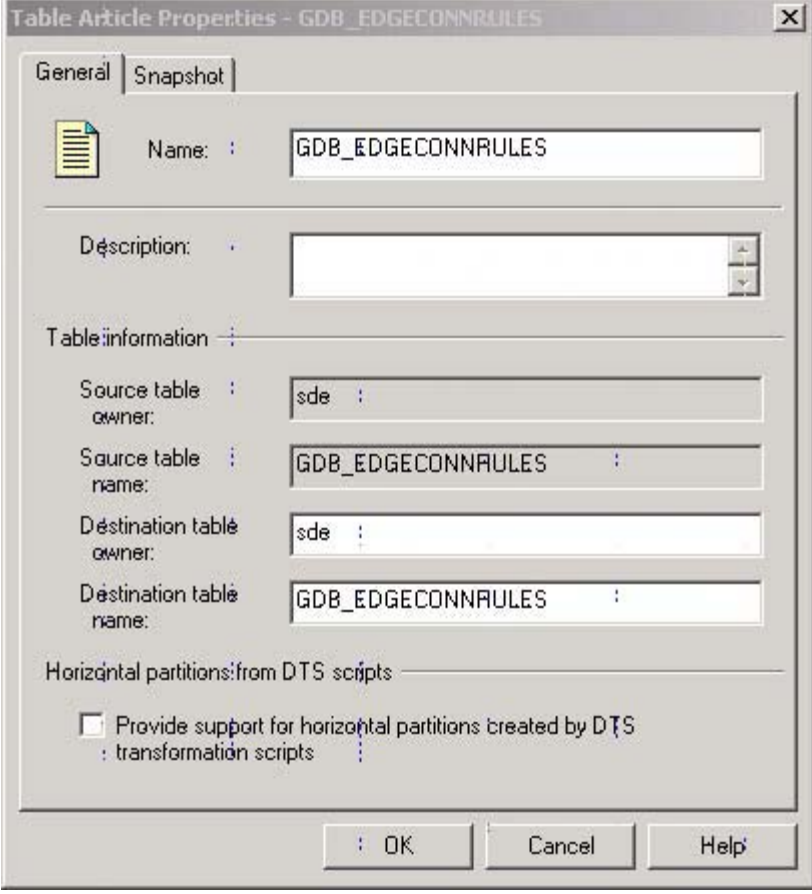
Select Snapshot publication, then click Next>.



For this example, SQL Server 2000 was run. ArcSDE currently supports SQL Server 2000 and SQL Server 7. Click Next>.



Choose the data to transfer; here all tables in sde will be replicated. It is necessary here to ensure that the tables are also replicated with the correct owner on the subscribing database. For each table that is replicated, click the ellipsis next to the Object and the following screen will appear. Enter the correct owner of the table for the subscribing database.



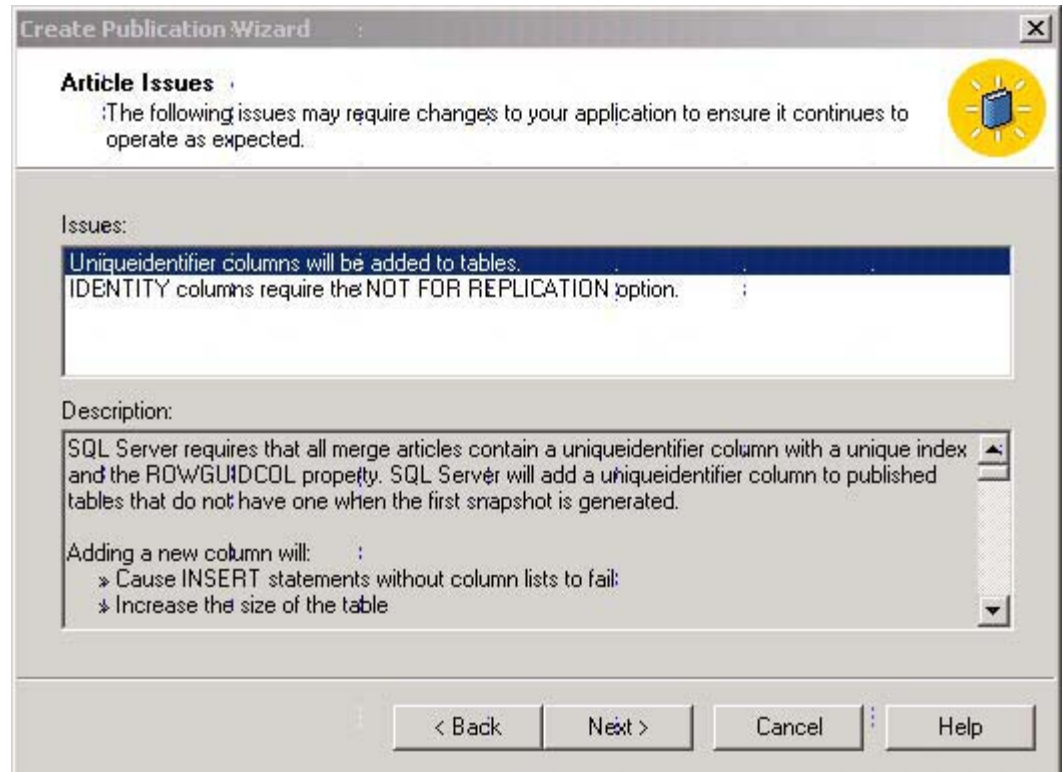
The image shows a Windows-style dialog box titled "Table Article Properties - GDB_EDGECONNRULES". It has two tabs: "General" (selected) and "Snapshot".

General Tab:

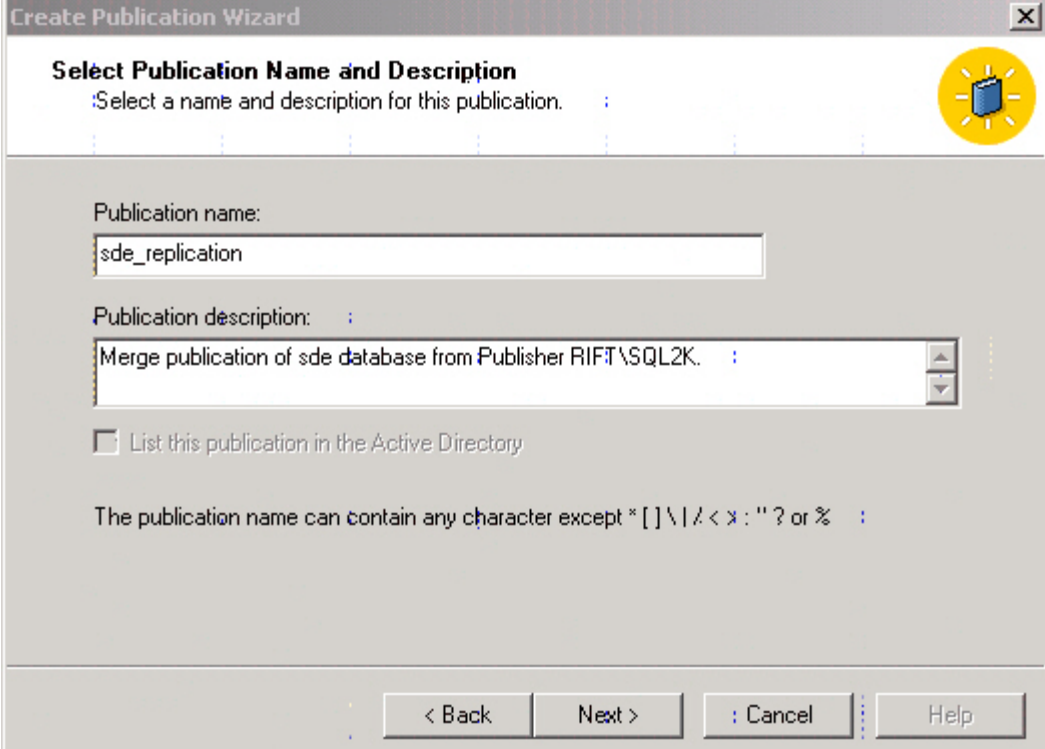
- Name:** GDB_EDGECONNRULES
- Description:** (empty text box)
- Table information:**
 - Source table owner:** sde
 - Source table name:** GDB_EDGECONNRULES
 - Destination table owner:** sde
 - Destination table name:** GDB_EDGECONNRULES
- Horizontal partitions:** from DTS scripts
 - ☐ Provide support for horizontal partitions created by DTS transformation scripts

At the bottom are three buttons: "OK", "Cancel", and "Help".

Confirm that the source table owner and the destination table owner are correct, then click OK to return to the previous screen. Click Next> at that screen to continue.



ROWGUIDCOL is a unique identifier that is added to tables used in merge replication. This makes it possible to track the lineage of changes made to a record. Click Next>.



The image shows a Windows-style dialog box titled "Create Publication Wizard". It has a standard title bar with a close button (X) in the top right corner. The main area is divided into sections. The first section is titled "Select Publication Name and Description" and contains the instruction "Select a name and description for this publication." To the right of this text is a yellow circular icon with a blue book and radiating lines. Below this, there are two text input fields. The first is labeled "Publication name:" and contains the text "sde_replication". The second is labeled "Publication description:" and contains the text "Merge publication of sde database from Publisher RIFT\SQL2K.". Below these fields is a checkbox labeled "List this publication in the Active Directory", which is currently unchecked. At the bottom of the dialog, there is a row of four buttons: "< Back", "Next >", "Cancel", and "Help". A status bar at the very bottom contains a small icon and some text.

Create Publication Wizard

Select Publication Name and Description
Select a name and description for this publication.

Publication name:
sde_replication

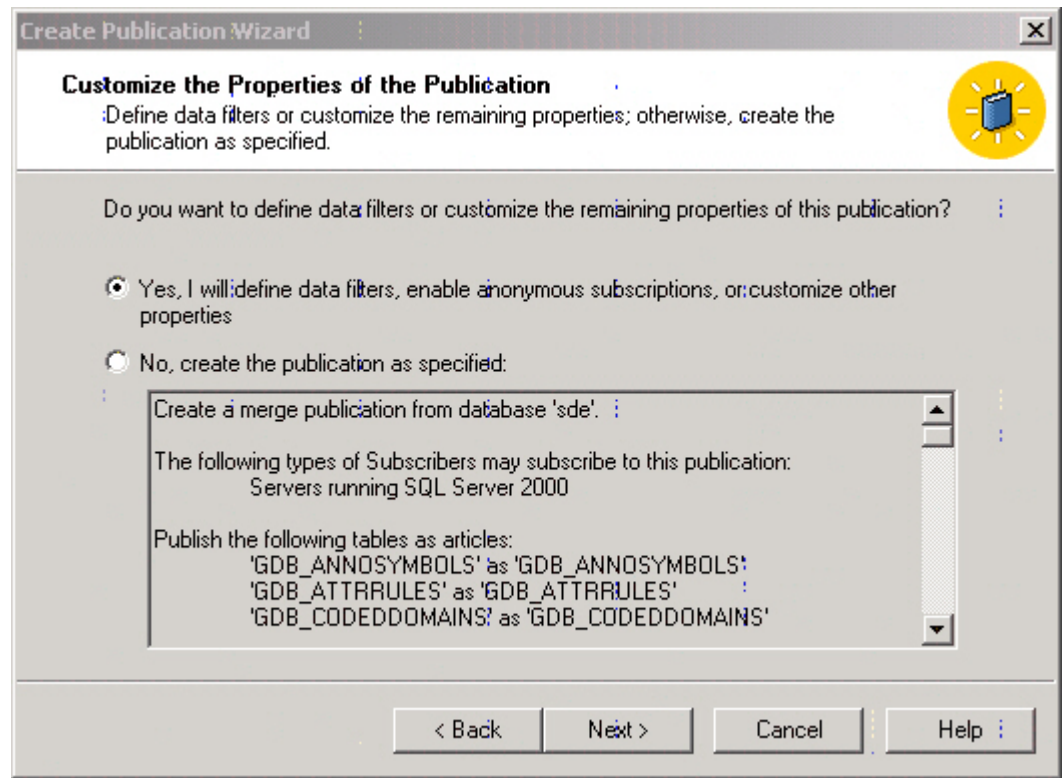
Publication description:
Merge publication of sde database from Publisher RIFT\SQL2K.

☐ List this publication in the Active Directory

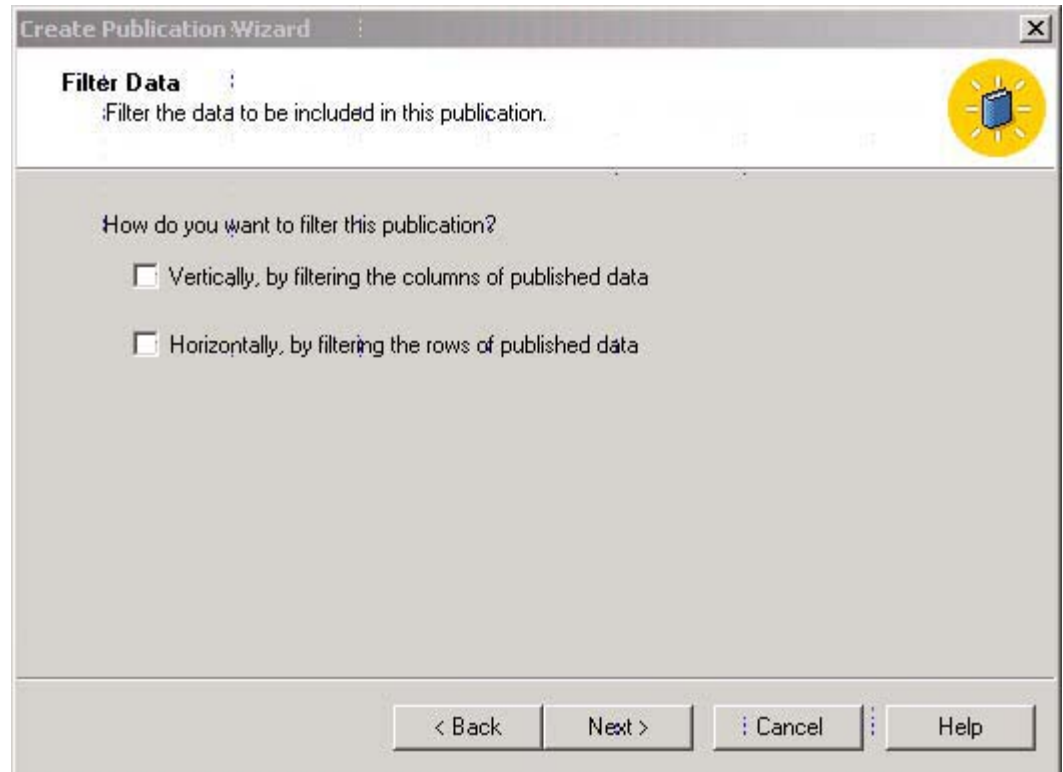
The publication name can contain any character except * [] \ / < > : " ? or %

< Back Next > Cancel Help

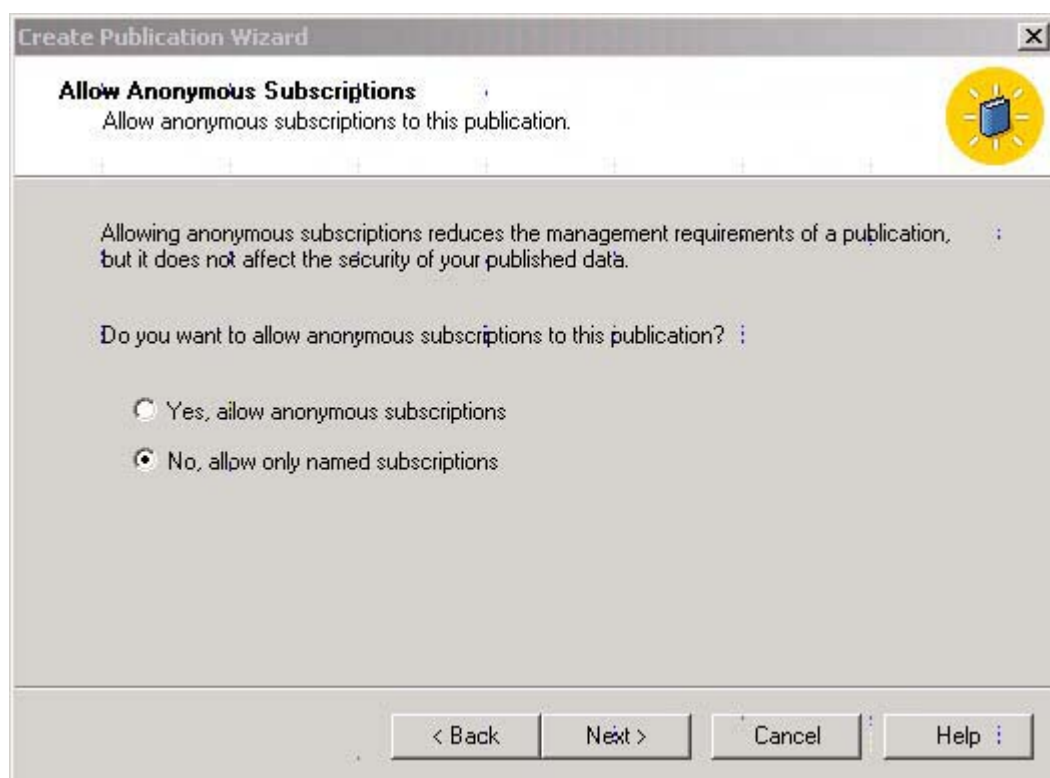
Enter a publication name and publication record and click Next>.



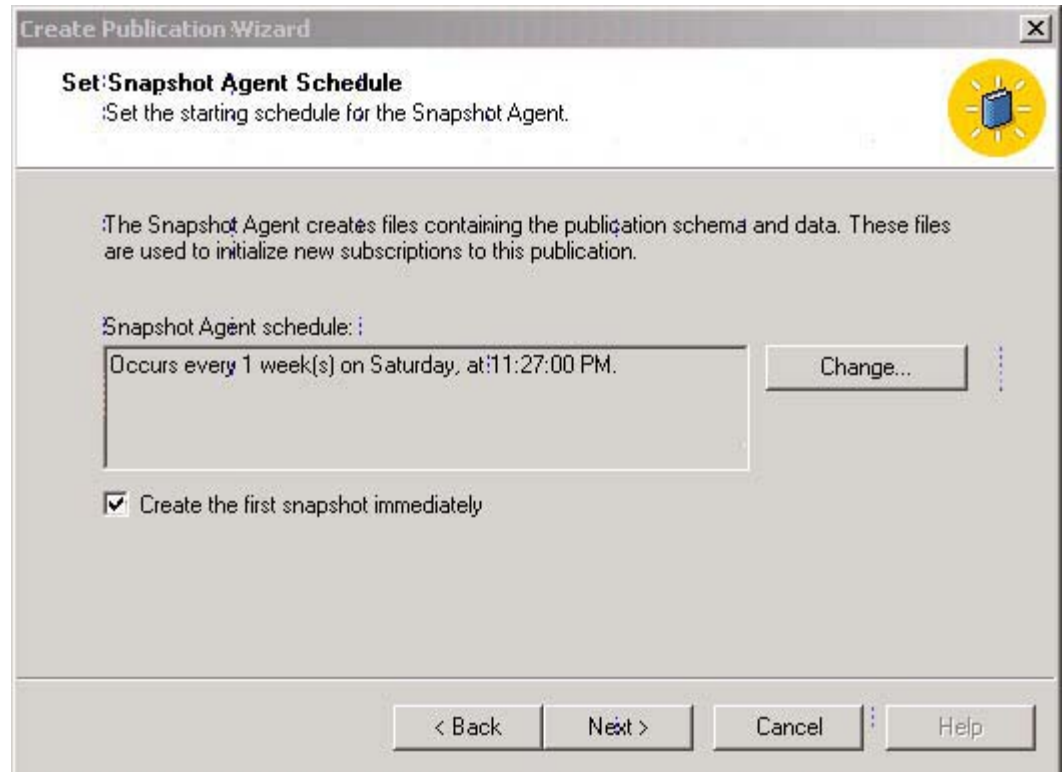
For this example, we will choose to customize the publication properties. Click Next>.



No filtering of this publication will be chosen. Click Next>.



Only named subscriptions will be allowed in this example. Click Next>.



The snapshot agent schedule will be changed. Click Change.

EdR Recurring Job Schedule - RIFT\SQL2K

Job name: (Initial Synchronization Schedule)

Occurs

☒ Daily
☐ Weekly
☐ Monthly

Daily

Every 1 day(s)

Daily frequency

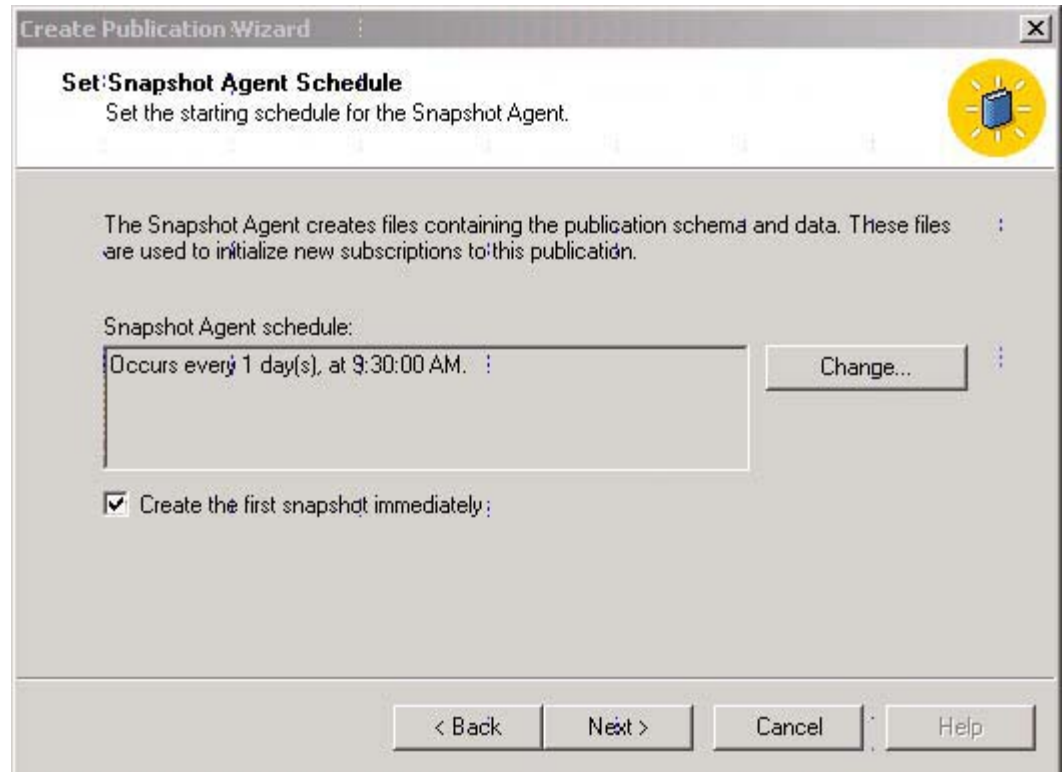
☒ Occurs once at: 9:20:00 AM
☐ Occurs every: 1 Hour(s) Starting at: 11:27:00 PM
Ending at: 11:59:59 PM

Duration

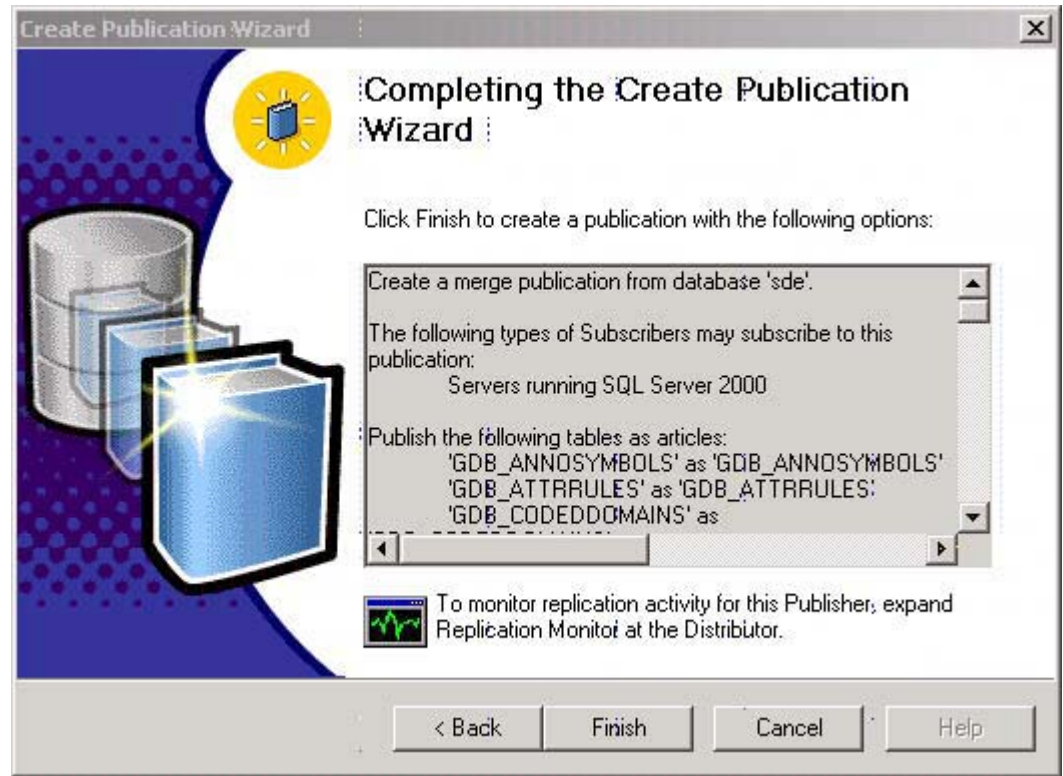
Start date: 12/13/2000 End date: 12/13/2000
☒ No end date

OK Cancel Help

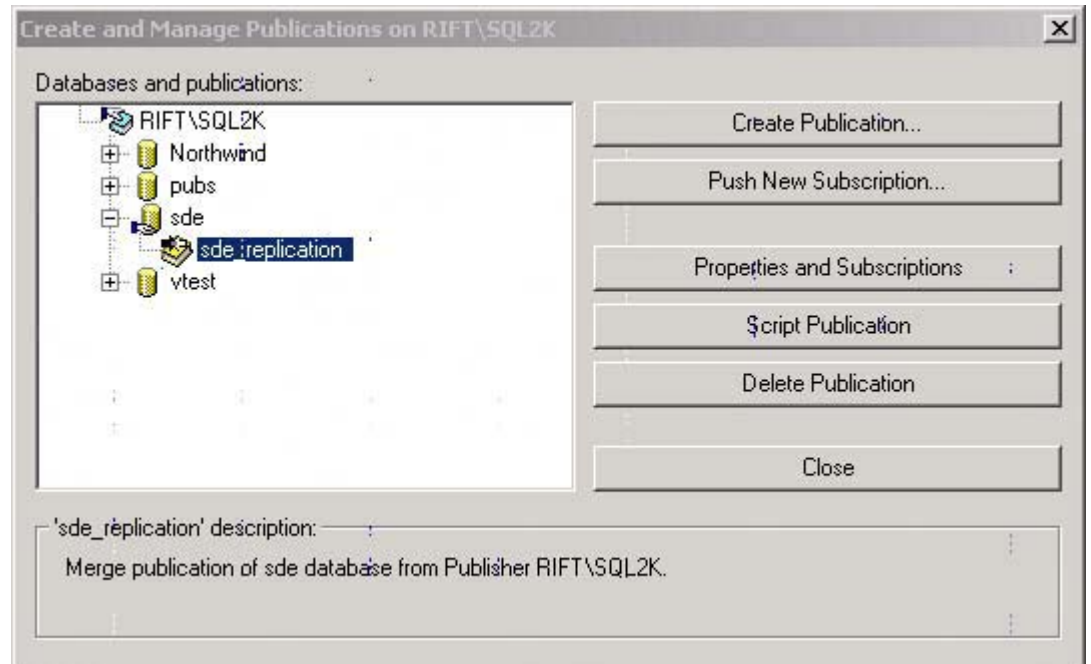
Set the parameters to make the snapshot. Click OK. You will return to the previous screen.



Click Next>.



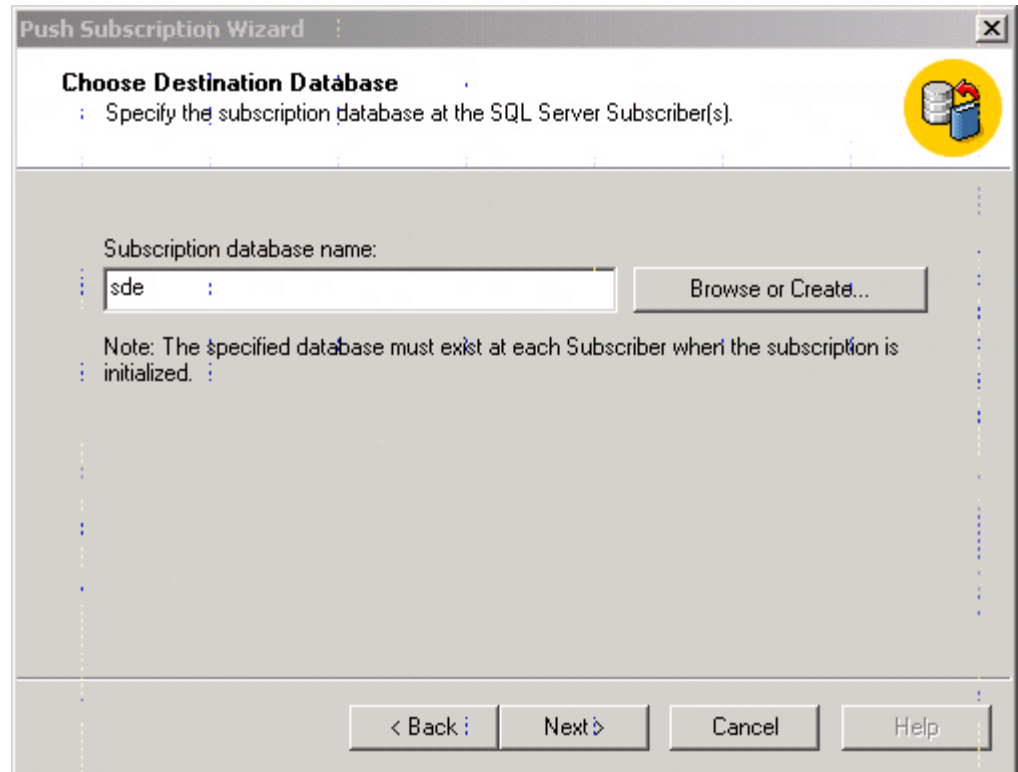
Click Finish to complete the creation of the publication. In this example we will create a push subscription of the publication.



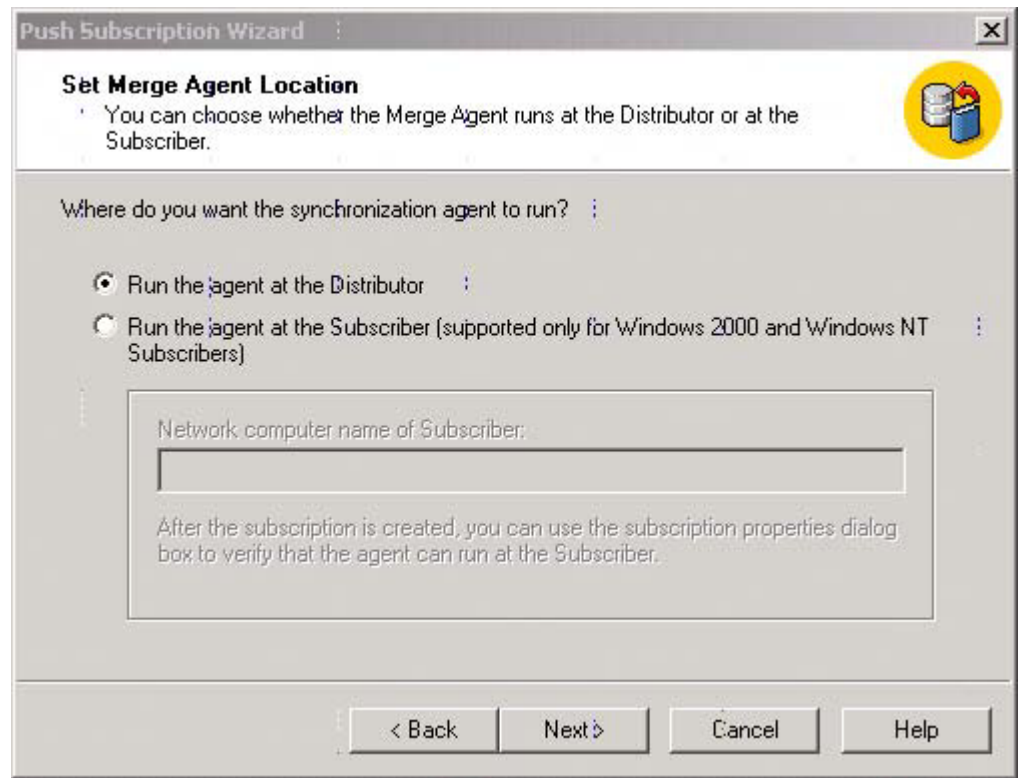
Highlight the publication and select Push New Subscription.



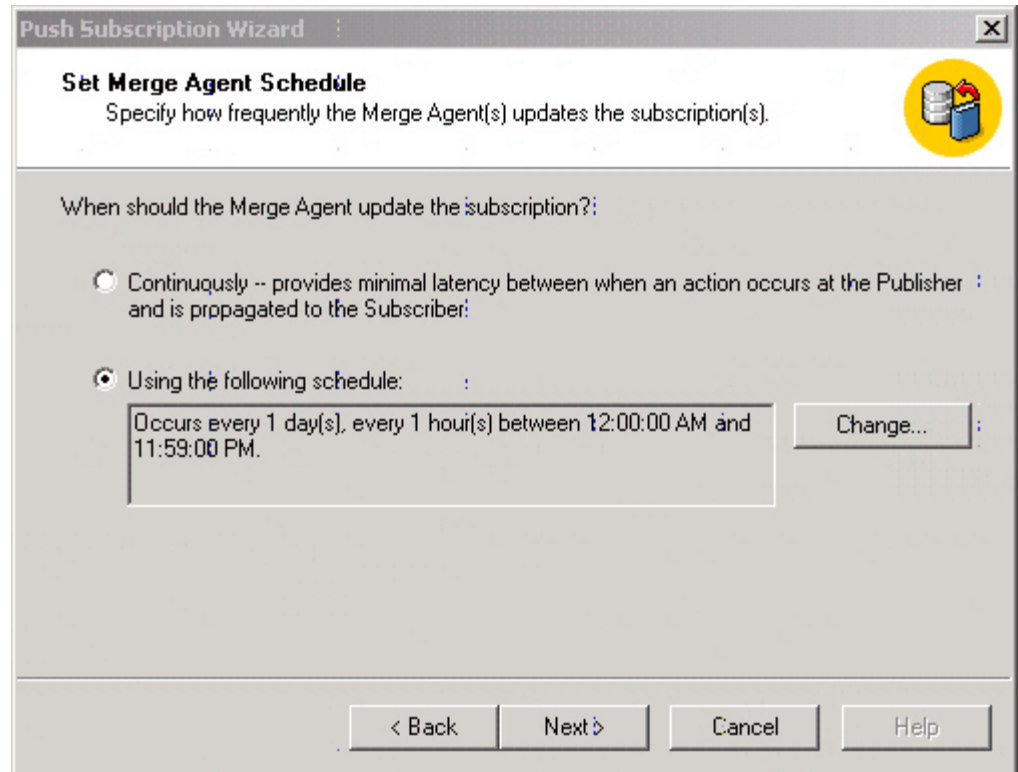
The Welcome to the Push Subscription Wizard will appear. This example will select advanced options. Click Next>.



Browse to or create the database on the subscribing database that will receive the publication. Click Next>.



For this example, the agent will be run at the Distributor. Click Next>.



This example will change the Merge Agent schedule rather than run continuously. Click Change.

Job name: (Default Distribution Agent Schedule)

Occurs

☒ Daily

☐ Weekly

☐ Monthly

Daily

Every 1 day(s)

Daily frequency

☒ Occurs once at: 10:00:00 PM

☐ Occurs every: 1 Hour(s) Starting at: 12:00:00 AM Ending at: 11:59:00 PM

Duration

Start date: 12/13/2008

End date: 12/14/2008

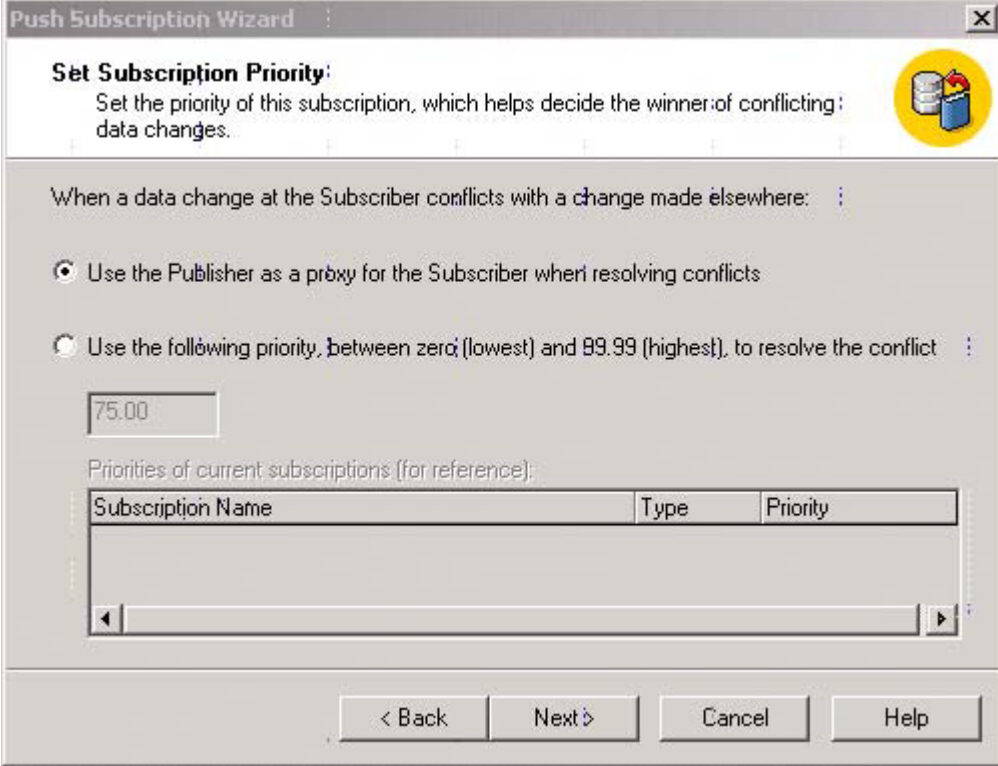
☒ No end date

OK Cancel Help

Select the parameters to run the Merge Agent schedule. Click OK.



For this example, choose Yes, initialize the schema and data. Click Start the Snapshot Agent to begin the initialization process immediately, based on the parameters to start earlier. Click Next>.



Push Subscription Wizard

Set Subscription Priority:
Set the priority of this subscription, which helps decide the winner of conflicting data changes.

When a data change at the Subscriber conflicts with a change made elsewhere:

- ☒ Use the Publisher as a proxy for the Subscriber when resolving conflicts
- ☐ Use the following priority, between zero (lowest) and 99.99 (highest), to resolve the conflict

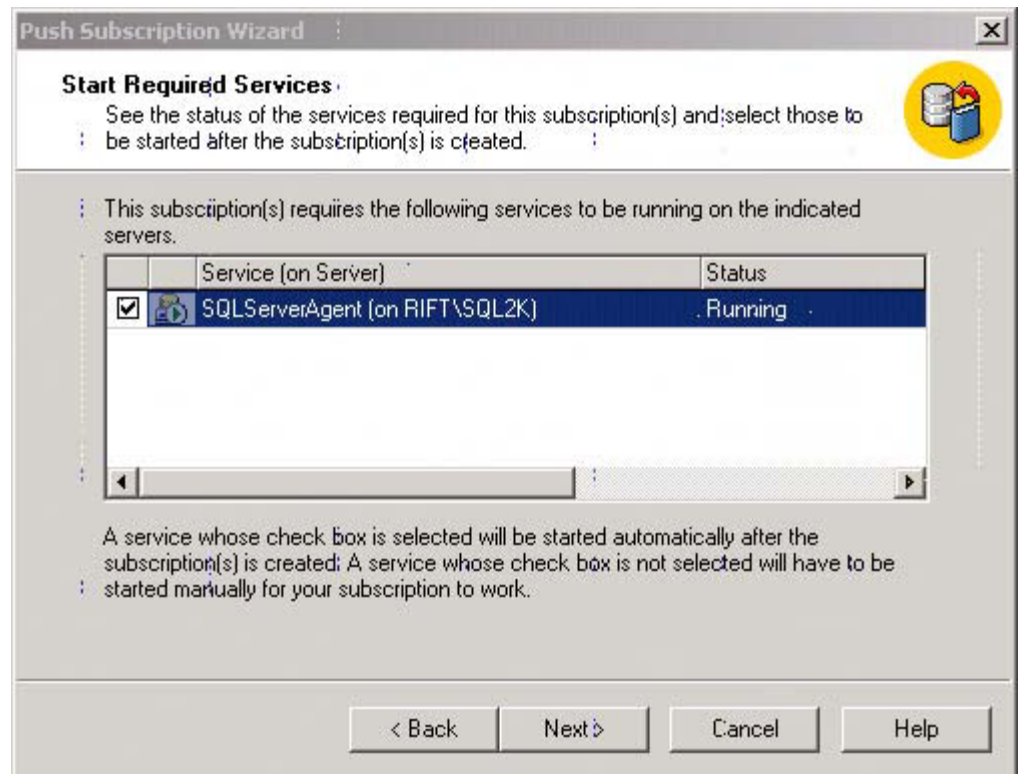
75.00

Priorities of current subscriptions (for reference):

Subscription Name	Type	Priority
-------------------	------	----------

< Back Next > Cancel Help

Select the subscription priority to resolve conflicts, then click Next>.



Click Next> to start the required services.



Click Finish to complete the replication procedure.

CHAPTER 10

Configuring transactional replication

ArcSDE 8.1 for MS SQL Server supports transactional and snapshot replication. The transactional model allows a replication distributor/publisher to update its data and then propagate those edits downstream to its subscribers. The snapshot model creates entire snapshots of data and delivers the entire snapshot to its subscribers. Snapshot replication works well with smaller, read-only data. Transactional replication proves better for frequent updates to dynamic data.

Before implementing transactional or snapshot replication, you should have a well-developed understanding of the SQL Server replication architecture. Make sure to read the section “Replication” in the SQL Server Books Online. This section demonstrates how to configure transactional and snapshot replication from a SQL Server 2000 publisher/distributor to a SQL Server 2000 replication subscriber.

Implementing transactional replication

Basic components of transactional replication with ArcSDE data

1. Create primary keys on all ArcSDE business tables. We recommend that you use a column that is an sde-maintained rowid. Use either the ObjectID field of data loaded with ArcCatalog or ArcToolbox or add a field with sdetable -o alter_reg and use the -C sde switch.
2. Create primary keys on five Geodatabase metadata tables (see the table below).
3. Define your publication: replicate all ArcSDE tables and stored procedures **except** SDE_process_information, SDE_state_locks, and SDE_table_locks. Also, it is not necessary to replicate the logfiles, SDE_logfiles and SDE_logfile_data. **Do not** replicate the execution of the SDE stored procedures. You must manually remap ownership of replication articles; SQL Server, by default, remaps all article ownership to ‘sa.’
4. Create Subscriptions to each replication publication.

5. Deliver your snapshot to your replication subscribers. You can do this by using the SQL Server Distribution Agent, backup and restore, or `sp_detach_db` and `sp_attach_db`. Using the Distribution Agent is easiest but not efficient for moving huge quantities of data across your network. Both backup and restore and `sp_detach_db` and `sp_attach_db` are also easy but will require you to run `sp_change_users_login` for each data owner and possibly regrant table permissions to all nonowners. Don't forget to copy the `SDE_process_information`, `SDE_state_locks`, and `SDE_table_locks` tables to the replication subscriber if you have used the Distribution Agent to move your data to the subscriber machines.
6. Connect to the Subscribers.

Step 1: Create primary keys

Transactional replication requires a primary key on each replication article. All ArcSDE tables except business tables will have a primary key already created. Some Geodatabase metadata tables do not have the primary key implemented; these will have to be created manually.

Create primary keys on your business tables

We recommend that you add a primary key to a column that has an sde-maintained row-id. Examples of such columns are OBJECTID columns created by loading data with ArcGIS tools ArcCatalog or ArcToolbox and columns set or created with `sdetable -o alter_reg`. Do not run `sdetable -o alter_reg` on data loaded with ArcCatalog/ArcToolbox for the purpose of creating a primary key. Use the existing objectid field instead. If you did not load your database with these tools but used `shp2sde`, `sdeimport`, etc., use the `alter_reg` command to create your field with the sde-maintained row-id column.

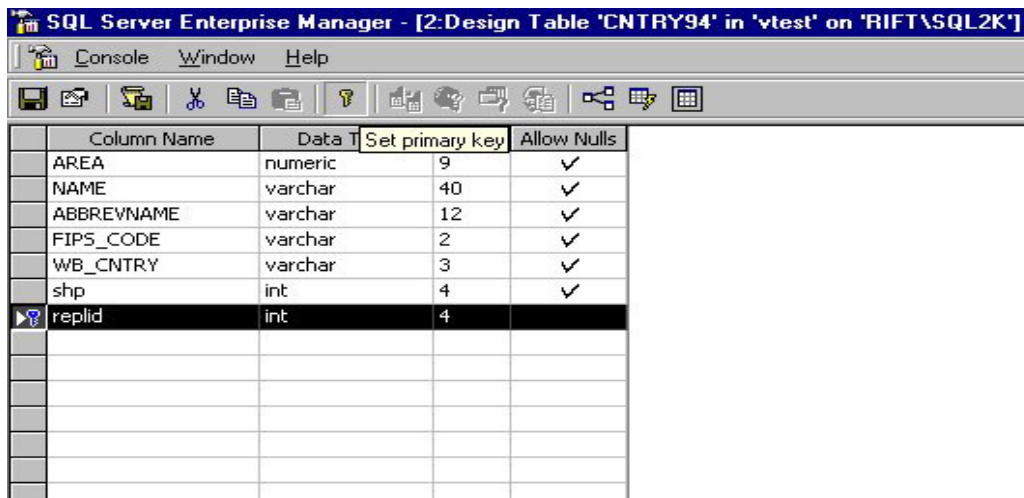
```
F:\world_sdex>sdetable -o alter_reg -i sql82k -t cntry94 -C sde -c replid
-u vtest -p go -D vtest
```

```
ArcSDE 8.1 Build 633 Tue Nov 21 22:30:10 PST 2000
Attribute Administration Utility
```

```
-----
Alter registration for table cntry94. Are you sure? (Y/N): y
Table cntry94's registration successfully altered.
```

To add the primary key constraint to this table, use the design table utility within the SQL Server Enterprise Manager. Find your business table, right-click it, and select “design table.” In the design table form, select the column on which you'll create the primary key. Now click the “Key” button.

Figure 10.1: In the Design Table dialog, select your primary key field and click the Key button.



Create primary keys on five Geodatabase metadata tables

This only applies to Geodatabase-dependent clients. If you are not using ArcGIS clients or did not load your data with ArcCatalog or ArcToolbox, skip this step. Otherwise, create a primary key on the following tables on the listed column(s).

Table	Column
GDB_FEATURECLASSES	ObjectClassID
GDB_NETWEIGHTASOCS	NetworkID
GDB_RELEASE	Any column
GDB_DEFAULTVALUES	ClassID,FieldName,Subtype
GDB_FIELDINFO	ClassID,FieldName

Step 2: Configure the Replication Publisher and Distributor

This step demonstrates using the SQL Server 2000 replication wizards. The process is similar in SQL Server 7. You can also use the replication stored procedures to set everything up. We recommend you use the wizards first to get familiar with the process, then either write out the stored procedures yourself or let the database script them for you.

Configure publishing and distribution

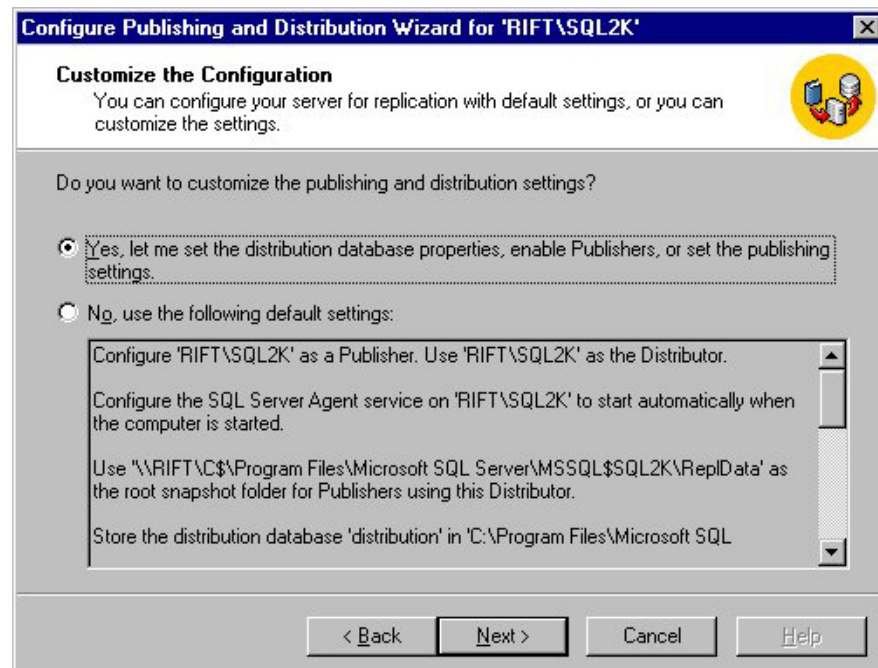
Start the **Configure Publishing and Distribution** Wizard by selecting and then right-clicking your server in the SQL Server Enterprise Manager. Choose View taskpad. This will repaint your details window with the SQL Server 2000 taskpad. Now click the Wizards tab. Under Setup Replication click Configure Publishing and Distribution

Wizard. You can click **Next** throughout this form as long as your server is listed in the entries for:

Specify '<server name>' or another server as a Distributor.
 Configure the properties of '<server name>' as a Distributor.
 Configure the properties of '<server name>' as a Publisher.

1. In the next form, **Select Distributor**, specify a distribution server if you don't want to take the default.
2. Click **Next** through the following form **Configure SQL Server Agent**. (NOTE: This agent may have to be owned by a domain account.)
3. Click **Next** through the following form **Specify Snapshot Folder**. (NOTE: You may get a warning message regarding sharing this folder.)
4. In the next form, **Customize the Configuration**, select "Yes, let me set the distribution database properties..." unless you are absolutely certain the defaults will work for you. In this case, we clicked Yes to customize settings regarding publication databases and enabling subscribers. Click **Next**.

Figure 10.2: Click Yes to customize the settings.



5. Click **Next** through the following form: Provide Distribution Database Information.
6. In the ensuing form, Enable Publishers, confirm that you have the correct publisher selected. In most cases, the default will be adequate. Click **Next**.
7. The next form, Enable Publication Databases, requires you to select which databases you want to enable for publication. In this case, we will enable two, the "sde" and "vtest" databases. You do this by selecting a type of replication per database. Since ArcSDE 8.1 does not support merge replication, you must only choose the "trans" option. Click **Next**.

Figure 10.3: Check the boxes in the trans column for each database you wish to enable for replication.

Configure Publishing and Distribution Wizard for RIFT\SQL2K

Enable Publication Databases
Enable databases for transactional (includes snapshot) or merge replication.

Databases:

Trans	Merge	Database Name
<input type="checkbox"/>	<input type="checkbox"/>	Northwind
<input type="checkbox"/>	<input type="checkbox"/>	pubs
<input checked="" type="checkbox"/>	<input type="checkbox"/>	sde
<input checked="" type="checkbox"/>	<input type="checkbox"/>	vtest

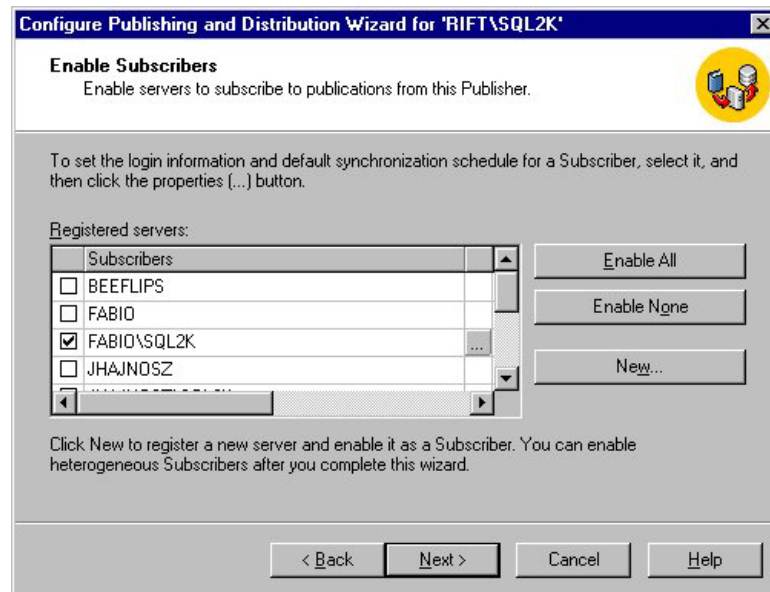
Transactional
Enable All
Enable None

Merge
Enable All
Enable None

A database owner can create a publication only after an administrator has enabled the database for publishing.

< Back Next > Cancel Help

8. Now you'll see the **Enable Subscribers** form. All the servers you have registered in your Enterprise Manager will be listed here. You may not want to enable them all for replication. In this case, we deselect all but one SQL Server 2000 instance (FABIO\sql2k) on a machine called "Fabio." Click **Next**.

Figure 10.4: Deselect all servers you **do not want** to enable for replication subscription.

You've arrived at the end of the Configure Publishing and Distribution Wizard. Congratulations! Review the values listed to make sure they are correct. Now click **Finish**. The configuration process will start, step through each point, and finish with this dialog if all went well.

Figure 10.5: Successful enabling of the distributor and publisher



Step 3: Create a publication

The MS SQL Server 2000 Books Online define a publication as:

A collection of one or more articles from one database. This grouping of multiple articles makes it easier to specify a logically related set of data and database objects that you want to replicate together.

Return to the taskpad view of your server and click the Wizards tab. Select the Create Publication topic under Setup Replication. In the ensuing form, Create and Manage Publications on <YourServerName>, you should see your databases listed in the tree view with the "hand" icon enabled to their left. Select the sde database and click the Create Publication button. The **Create Publication** Wizard will start. Click **Next**.

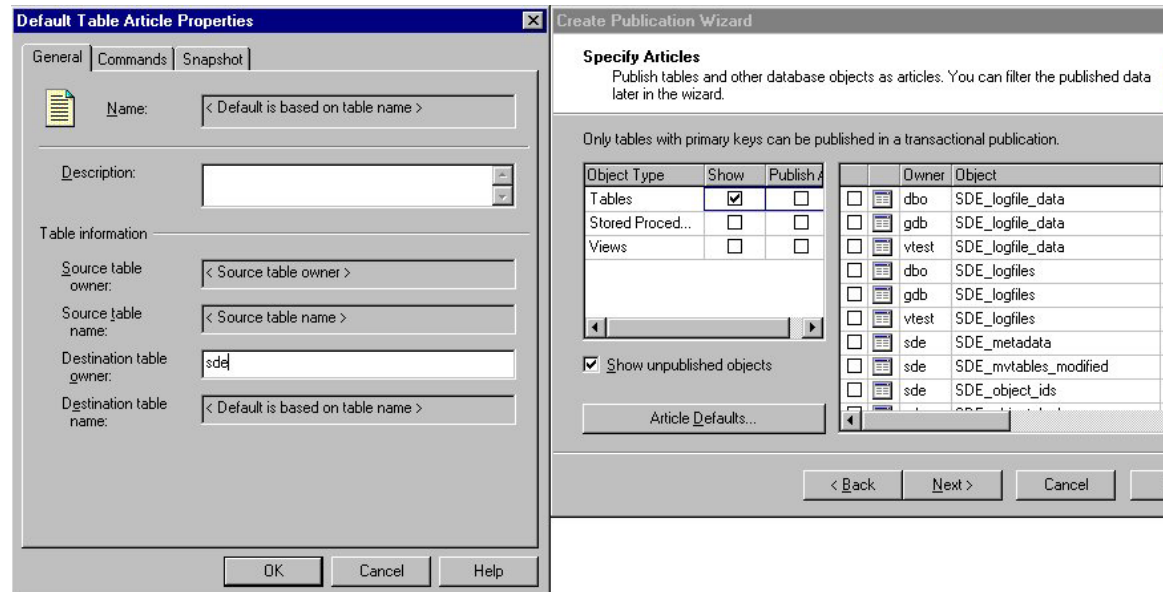
1. The **Choose Publication Database** wizard appears. You should see your sde database highlighted in this box. Click **Next**.
2. The **Select Publication Type** dialog is next. Click the Transactional Replication radio button. Click **Next**.
3. Now the **Specify Subscriber Types** form displays. In this case, you have enabled one other server, a SQL Server 2000 machine, for transactional replication. Select all your platforms that apply here. You can replicate to SQL Server 2000 and SQL Server 7.0 or heterogeneous data sources such as DB2 or MS Access. Click **Next**.
4. You should now see the **Specify Articles** form. A replication “article” is any table, stored procedure, or view you wish to replicate. In this form, you’ll have to specify which articles you want to replicate and who should own those articles on the subscriber servers.

NOTE: You must ensure that your articles are correctly mapped to their proper owners. SQL Server, by default, will change all table, view, and stored procedure ownership to ‘sa.’ This will prevent the sde service from starting on the subscriber machine or client direct connections from connecting to the data store.

There are several ways to remap the article permissions. This example describes one method but assumes this condition: all replicated tables, stored procedures, and views in the ‘sde’ database are owned by the sde user, and everything in the ‘vtest’ database is owned by user ‘vtest.’ This condition allows the use of the **Article Defaults** dialog in the SQL Server 2000 **Specify Articles** form.

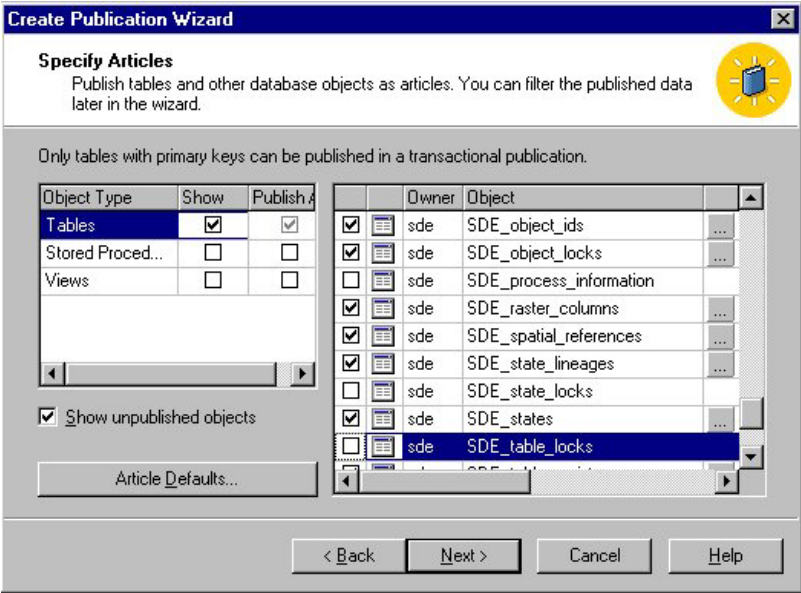
With focus on the Tables row of the Object Type column in the grid control on the left of this form, click the **Article Defaults** button. Select Table Defaults, OK. The **Default Table Article Properties** form shows. For Destination table owner input ‘sde.’ Click **OK** to return to the **Specify Articles** form.

Figure 10.6: Click the Article Defaults button to open the Default Table Article Properties Form. Input ‘sde’ into the Destination table owner field.



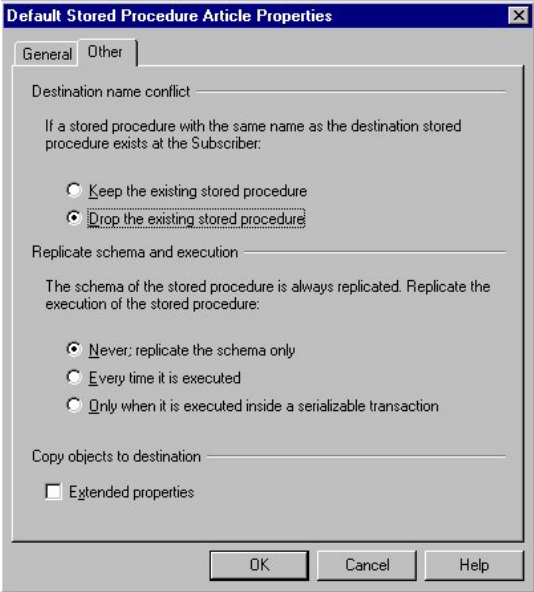
Back in the **Specify Articles** form, click the Publish All check box in the Tables Object Type record in the grid control on the left. In the right-hand grid, deselect all the logfile tables, the SDE_process_information, SDE_table_locks, and SDE_object_locks tables.

Figure 10.7: Deselect all logfile tables, the SDE_process_information, SDE_table_locks, and SDE_object_locks tables.



Now shift focus to the Stored Procedures record in the **Specify Articles** form by clicking it. Click the **Article Defaults** button to show the **Default Stored Procedure Article Default Properties** form and input 'sde' into the owner text field. Click the "Other" tab. Ensure that this screen has the "Never; replicate the schema only" radio button checked (see Figure 10.8).

Figure 10.8: Ensure that the "Never; replicate the schema only" radio button is checked.



Now perform step B for the stored procedures, only do not deselect anything.

Redo steps A and B for views, only do not deselect anything.

Check a few of the stored procedures and tables to ensure they have the proper ownership by clicking the small button on the right

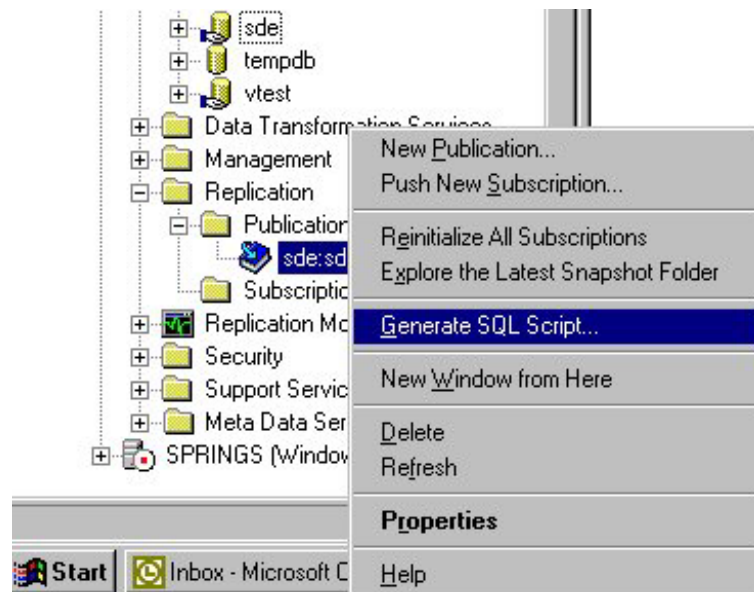
Click **Next**.

5. Click **Next** all the way to the end of the wizard.

Shortcut for point 4:

The steps in point 4 may seem a bit cumbersome. If you feel comfortable with publications in general, take all the defaults. This will remap all of your table, stored procedure, and view ownership to sa. When you've created the publication, script it to SQL by right-clicking the publication in the Enterprise Manager and select Generate SQL Scripts (see Figure 10.9).

Figure 10.9: Generate SQL scripts for your publication.



Save the SQL Script and delete your publication. Open the script in the SQL Server Query Analyzer and replace all instances of the string `@destination_owner = N'<some owner Name will be here>'` with the proper owner. For example:

```
...@destination_owner = N'terra\slim'
...@destination_owner = N'sde'
```

Now run the entire script to re-create the publication with the proper ownership of your replicated articles.

Create a publication for each database to be replicated. This means you'll have to rerun steps 1–5 for each database you want to replicate.

Step 4: Create subscriptions

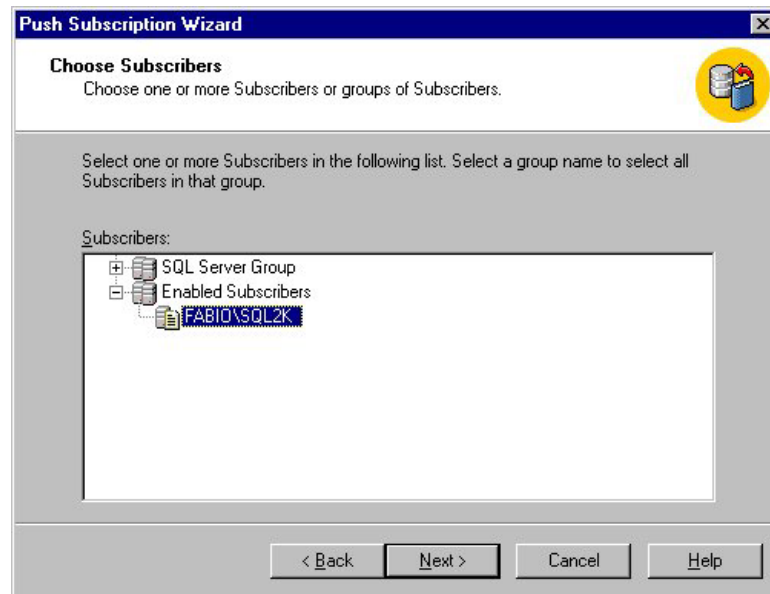
The SQL Server Books Online lists a Subscriber as:

Servers that receive replicated data. Subscribers subscribe to publications, not to individual articles within a publication.

In this step, we will create a new push subscription on a SQL Server 2000 instance named “Fabio\sql2k.” Start the **New Subscription Wizard** by right-clicking the publication in the Enterprise Manager’s Replication branch publication node. Select “Push New Subscription.” Click **Next** in the Welcome screen.

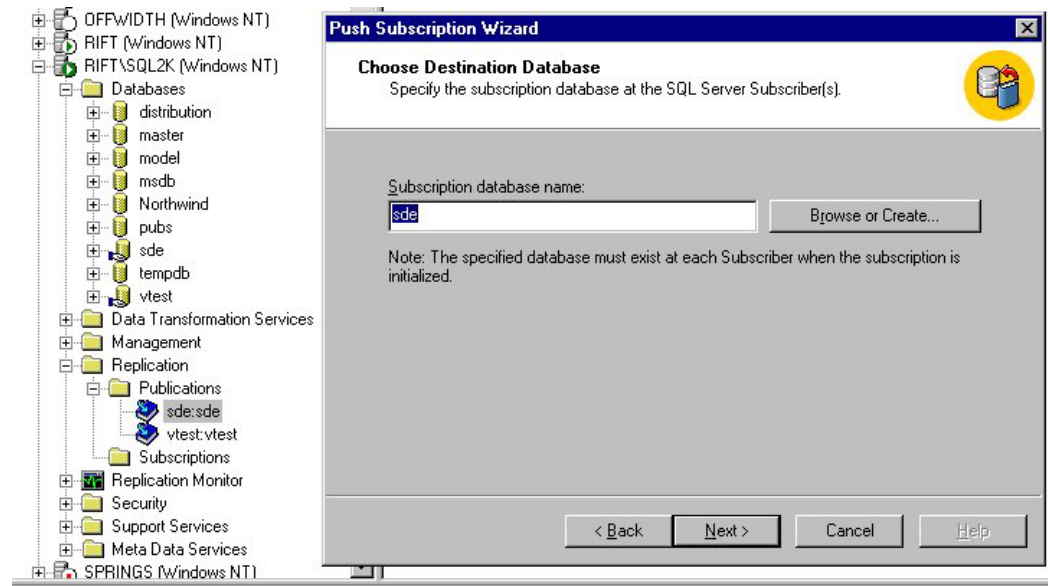
The Choose Subscribers form appears. You should see at least one enabled subscriber listed here. If you don’t, you must go back and enable a subscriber for replication. Select one or more of the listed subscribers. See Figure 10.10.

Figure 10.10: You should see a list of enabled subscribers. Select one or more.



Click **Next** to proceed to the **Choose Destination Database** form. You’ll see a text box with a selected database name here. If you right-clicked the ‘sde:sde’ publication, then you should see the word ‘sde’ written here. See Figure 10.11.

Figure 10.11: Detail of **Choose Destination Database** form. Note the selected publication in the Enterprise Manager 'sde:sde.'



Click **Next>**. Now the **Set Distribution Agent Schedule** appears. This form controls latency between the publisher and the subscribers. If you must maintain close synchronization between the publisher and subscribers, then take the default, continuous updates. You can set the schedule to update as you please. In this case, we take the defaults. Click **Next>**.

The **Initialize Subscription** form opens. This form controls the subscription initialization on the subscriber. If you want the Snapshot and Distribution Agents to initialize your schema and deliver your data to the subscriber machines, then take the defaults. If you've moved or will move your data to the subscriber yourself, click the radio button labeled "No, the Subscriber already has the schema and data." In this example, take the default and check the "Start the Snapshot Agent to begin the initialization process immediately" check box. This means that the Snapshot and Distribution Agents will initialize the schema and move all the articles to the subscriber.

Figure 10.12: Check the “Start the Snapshot Agent.” check box.



Click **Next>**.

The **Start Required Services** Form displays. Make sure the status of your SQL Server Agent is “Running.” Click **Next>**.

You’ve reached the last screen in the **Push Subscription Wizard**. Click **Finish** to start the agents.

Do points 1–7 for each publication you have. In this case, we will have to redo 1 to 7 for the ‘vtest’ publication.

Step 5: Deliver your snapshot to your replication subscribers

If you have followed these steps exactly, then most of your data has been moved to the subscribers. However, you may find that the Agents are too slow in initializing your schema and moving your data. You can do this manually by backing up your data to be replicated and restoring it to the subscriber machine. See point 4 of Step 4 for more information on what to do in the **Create Publication Wizard**. The next section explains how to manually distribute your replication snapshot.

Manually initializing your replication snapshot

Again, if you have followed these instructions verbatim, you don’t have to do this. This only applies to scenarios that want to distribute data manually, outside the scope of the distribution and snapshot agents.

Move your data to the subscriber machine. You can do this with backup and restore, `sp_detach_db` and `sp_attach_db`, or Data Transformation Services (DTS).

If you’ve ran backup and restore or `sp_detach_db` and `sp_attach_db`, you must run `sp_change_users_login` to synchronize your SQL Server logins and users.

Moving the SDE_process_information, SDE_state_locks, and SDE_table_locks tables to the subscribers

You must move three ArcSDE tables to the replication subscriber, the SDE_process_information, SDE_state_locks, and SDE_table_locks tables. The easiest way to do this is to use the 'DTS.' Make sure that these tables, when moved to the subscribers, are still owned by the 'sde' user.

Special considerations for MS SQL Server 7.0

SQL Server 7.0 and 2000 have some differences with regard to replication. The most important one has to do with replicating stored procedures. With SQL Server 2000, you can replicate the schema of the stored procedures but not their execution. This option is not available with SQL Server 7.0. SQL Server 7.0 will replicate the stored procedures to the subscribers only if you allow their execution. This option is not supported in ArcSDE 8.1. **For SQL Server 7.0 replication of ArcSDE 8.1 data, you'll have to move all the stored procedures to each subscriber manually.**

Step 6: Connect to the subscribers

Now you should be able to make connections to the subscriber machines serving ArcSDE data. You can either start an SDE service on each subscriber or make direct connections from ArcSDE clients such as ArcGIS or both. If you want to start an ArcSDE service on each subscriber, then you must have the ArcSDE application installed on those servers. However, you don't have to install the ArcSDE application on a server setup for Direct Connections. See Chapter 5, 'Connecting to SQL Server', in this manual for more information on this.

APPENDIX A

ArcSDE compressed binary

This chapter describes the state of ArcSDE with no configuration changes made to the server or any spatial databases. This chapter also details how the ArcSDE application manages your data with regard to indexes, triggers, and specific SQL Server configuration settings. This chapter will help you understand how to manage spatial data.

Compressed binary

After the client verifies the geometry, it compresses it and sends it to the server, where it is stored in compressed binary format in a feature table. Compressing the geometry on the client offloads the task from the ArcSDE server and reduces the transmission time to send the geometry to the ArcSDE server. Storing compressed geometry data reduces the space required to store data by as much as 40 percent.

Compressed binary metadata tables

A compressed binary feature class comprises three tables: the business table, the feature table, and the spatial index table.

The business table contains attributes and a spatial column. The spatial column is a key to the feature and spatial index tables.

The relationship between the business table and the feature table is managed through the spatial column and the FID column. This key, which is maintained by ArcSDE, is unique.

NAME	DATA TYPE	Allow Nulls
fid	INTEGER(4)	NOT NULL
numofpts	INTEGER(4)	NOT NULL
entity	SMALLINT(2)	NOT NULL
eminx	FLOAT(8)	NOT NULL
eminy	FLOAT(8)	NOT NULL
emaxx	FLOAT(8)	NOT NULL
emaxy	FLOAT(8)	NOT NULL
eminz	FLOAT(8)	NULL
emaxz	FLOAT(8)	NULL
min_measure	FLOAT(8)	NULL
max_measure	FLOAT(8)	NULL
area	FLOAT(8)	NOT NULL
len	FLOAT(8)	NOT NULL
points	IMAGE	NULL
anno_text	VARCHAR(255)	NULL

Feature table schema

The feature table stores the geometry, annotation, and CAD in the image type POINTS column. The internal SQL Server datatype is listed in the table above. The ArcSDE datatype for each column is defined below.

- ✎ fid (SE_INTEGER_TYPE)—contains the unique ID that joins the feature table to the business table
- ✎ entity (SE_INTEGER_TYPE)—the type of geometric feature stored in the shape column (e.g., point, linestring)
- ✎ numofpts (SE_INTEGER_TYPE)—the number of points defining the shape
- ✎ eminx, eminy, emaxx, emaxy (SE_FLOAT_TYPE)—the envelope of the shape
- ✎ eminz (SE_FLOAT_TYPE)—the minimum Z value in the shape
- ✎ emaxz (SE_FLOAT_TYPE)—the maximum Z value in the shape
- ✎ min_measure (SE_FLOAT_TYPE)—the minimum measure value in the shape
- ✎ max_measure (SE_FLOAT_TYPE)—the maximum measure value in the shape
- ✎ area (SE_FLOAT_TYPE)—the area of the shape
- ✎ len (SE_FLOAT_TYPE)—the length or perimeter of the shape
- ✎ points (SE_SHAPE_TYPE)—contains the byte stream of point coordinates that define the shape's geometry
- ✎ anno_text (SE_STRING_TYPE)—contains the feature annotation string.

NAME	DATA TYPE	Allow Null?
sp_fid	INTEGER(4)	NOT NULL
Gx	INTEGER(4)	NOT NULL
Gy	INTEGER(4)	NOT NULL
Eminx	INTEGER(4)	NOT NULL
Eminy	INTEGER(4)	NOT NULL
Emaxx	INTEGER(4)	NOT NULL
Emaxy	INTEGER(4)	NOT NULL

Spatial index table schema

The spatial index table defines the grid range and extent of all geometry in an ArcSDE feature class.

- sp_fid (SE_INTEGER_TYPE)—contains the unique ID that joins the feature table to the business table
- gx/gy (SE_INTEGER_TYPE)—defines the feature's extent in grid cells
- eminx/eminy/emaxx/emaxy (SE_INTEGER_TYPE)—defines the extent of the feature in system units

In this example the FEATURE-ID column from the WELLS business table references features from the feature and spatial index tables:

WELL_ID	DEPTH	ACTIVE	FEATURE-ID
1	30029	Yes	101
2	13939	No	102
3	92891	No	103
...

FID	AREA	LEN	EMINX,EMINY,...	POINTS
101				<compressed feature>
102				<compressed feature>
103				<compressed feature>
...				...

SP_FID	GX	GY	{EMINX,EMINY,EMAXX,EMAXY}
101	70	100	
102	70	100	
103	71	100	
...			

A business/feature/spatial index key reference

The spatial grid index

The spatial grid index is a two-dimensional index that spans a feature class such as the reference grid you might find on a common road map. You may assign the spatial grid index one, two, or three grid levels, each with its own distinct cell size. The mandatory first grid level has the smallest cell size. The optional second and third grid cell levels are disabled by setting them to 0. If enabled, the second grid cell size must be at three times larger than the first grid cell size, and the third grid cell size must be three times larger than the second grid cell size.

The spatial index table (S<feature class_id>) has seven integer columns that store the grid cell values, feature envelopes, and corresponding feature IDs. Adding a feature to a feature class adds one or more grid cells to the spatial index table. The number of records added to the spatial index table depends on the number of grid cells the feature spans.

The spatial index table contains two indexes. One index is on the SP_FID column, which contains the feature ID. The other is a clustered composite index that includes all of the columns of the spatial index table. Since all of the columns of the spatial index table are indexed, the values of the table are read from the leaf blocks of the index and not the table data blocks. The result is less I/O and better performance. In addition, the spatial index table is not accessed whenever the feature class is queried. Therefore, when considering how to position the tables and indexes to reduce disk I/O contention, you should be concerned about the positioning of the indexes of the spatial index table but not the table itself.

Building the spatial index

Every time a feature class is added to a business table, a persistent spatial index is built for it.

The ArcSDE server manages the spatial index throughout the life of the feature class. As features are inserted, updated, or deleted, the spatial index is automatically updated.

A load-only mode disables spatial index management until loading is completed. This boosts loading performance substantially and is imperative for bulk loading efforts (no queries are allowed in the load-only mode except native SQL-based queries).

Once loading has been completed, the spatial index is enabled by returning it to normal I/O mode. The conversion from normal I/O mode to load-only I/O mode reconstructs the spatial index.


Inserting, updating, or deleting a feature updates the spatial index when the feature class is in normal I/O mode.

ArcSDE overlays the extent of each feature onto the lowest grid level to obtain the number of grid cells. If the feature exceeds four cells, ArcSDE promotes the feature to the next highest grid level, if you have defined one. ArcSDE will continue to promote the feature until it fits within four cells or less or until the highest defined grid level is reached. On the highest defined grid level, geometries can be indexed by more than four grid cells.

ArcSDE adds the feature's grid cells to the spatial index table with their corresponding shape ID and feature envelope. The grid level is encoded with each grid cell.

In the example below, the feature class has two grid levels. Area shape 101 is located in grid cell 4 on level 1. A record is added to the spatial index table because the feature resides within four grid cells (in this case it is one). Area feature 102's envelope is located in cells 1 through 8 on level 1. Because the feature's envelope resides in more than four grid cells, the feature is promoted to level 2, where its envelope fits within two grid cells. Feature 102 is indexed at level 2, and two records are added to the spatial index table.

Figure A.1: Shape 101 is indexed on grid level 1, while shape 102 is indexed on grid level 2 where it is in only two grid cells.



Spatial queries and the spatial index

Spatial queries, like finding all the lakes within a state boundary, use the spatial index. A spatial index is used unless the search order has been set to `SE_ATTRIBUTE_FIRST` in the `SE_stream_set_spatial_constraints` function. When the search order is set to `SE_ATTRIBUTE_FIRST`, ArcSDE ignores the spatial index, and the criteria of the attribute where clause determines which records of the feature class the query returns.

Whenever the spatial index is used, the ArcSDE service generally uses the following decision process when performing the query:

1. Define the envelope. The envelope could be defined directly by the application such as the extent defined by the ArcMap zoom in tool. Alternatively, the envelope may be defined as the envelope of another feature.
2. Join the spatial index table with the feature table and return all features whose grid cells intersect the envelope.
3. Join the feature table with the business table and apply the criteria of the attribute where clause to further refine features returned.

Grid cell size impacts the size of the spatial index table. Setting up the spatial index means balancing the cell sizes—smaller cell sizes mean more cells per shape, which requires more entries in the spatial index table.

Guidelines for tuning the spatial index

Because client applications and spatial data profiles vary from one system to another, no single solution fits all. Experienced users of ArcSDE often experiment with the spatial index, trying different cell sizes and different grid level configurations.

The `sdelayer` command has several operations that can help you optimize the spatial index by changing the grid cell sizes and adding new grid levels with the 'alter' operation. The 'stats' and 'si_stats' operations profile your spatial data and current spatial index.

The following guidelines can help improve the performance of spatial queries.

- € Consider how many grid levels are needed and remember that the ArcSDE server scans the spatial index table once per grid level. Often a single grid level

is the best solution for a feature class, despite the notion of distributing geometries evenly across many grid levels to minimize the spatial index entries.

- € Use one grid level for pure point type feature class and consider making the cell sizes large. Spatial queries generally process point geometries faster than other geometry types.
- € Monitor the spatial index. Tuning a spatial index is difficult if the data changes frequently. Tuning depends on the structure of the spatial data. Periodically assess the spatial index as your spatial data changes.
- € Base the spatial index on the application. Match the spatial index grid cell sizes to the extent of the application window. By doing so, the application is probably viewing exact entries in the spatial index table. This helps to size the spatial index table suitably and reduces the amount of processing because fewer candidate feature IDs must be evaluated against the feature table (see ‘Spatial queries and the spatial index’ above).
- € For unknown or variable application windows, start by defining one grid level with a cell size three times the average feature extent size. Query the feature table to obtain the average feature size with the following SQL statement:

```
select (avg(emaxx - eminx) + avg(emaxy - eminy)) / 2
from f<N>;
```

(where <N> is the layer number of the feature class)

Such spatial index configuration minimizes the number of rows in a spatial index table while maintaining the proficiency of the index because the majority of features can be referenced by less than one or two grid cells.

- € Which application loaded your data? The dataloaders used in ArcCatalog and ArcToolbox will not calculate an optimal spatial index. The ArcSDE command tool shp2sde will generally calculate a better grid size. An easy test is to load the same data with both toolsets (e.g., ArcCatalog and shp2sde) and compare calculated grid sizes and drawing times for each. Alternatively, load the data with ArcCatalog or ArcToolbox and run the query listed immediately above. Use sdelay -o alter to change your grid size.
- € Design the feature class around spatial data categories such as type, geometry size, and distribution. Sometimes a carefully designed feature class, using these categories, can substantially boost the performance of spatial queries.

Displaying spatial index statistics

The sdelay command’s spatial indexstatistics operation, ‘si_stats’, can help you determine optimum spatial index grid sizes. Optimum grid cell sizes depend on the spatial extent of all feature geometries, the variation in feature geometry spatial extent, and the types of searches to be performed on the map feature class. Below is a sample output generated by si_stats:

```
$ sdelay -o si_stats -l victoria,parcels -u av -p mo -i sde81
ArcSDE 8.1 Wed Jan 17 22:43:09 PST 2000
Layer Administration Utility
-----
Layer 1 Spatial Index Statistics:
Level 1, Grid Size 200
|-----|
| Grid Records: 978341 |
| Feature Records: 627392 |
| Grids/Feature Ratio: 1.56 |
| Avg. Features per Grid: 18.26 |
| Max. Features per Grid: 166 |
| % of Features Wholly Inside 1 Grid: 59.71 |
```

```

-----
| Spatial Index Record Count By Group |
| Grids:  <=4  >4  >10  >25  >50  >100  >250  >500 |
|-----|
| Shapes: 627392  0  0  0  0  0  0  0 |
| % Total: 100%  0%  0%  0%  0%  0%  0%  0% |
|-----|
Level 2,  Grid Size 1600 (Meters)
|-----|
| Grid Records: 70532 |
| Feature Records: 36434 |
| Grids/Feature Ratio: 1.94 |
| Avg. Features per Grid: 18.21 |
| Max. Features per Grid: 82 |
| % of Features Wholly Inside 1 Grid: 45.35 |
|-----|
| Spatial Index Record Count By Group |
| Grids:  <=4  >4  >10  >25  >50  >100  >250  >500 |
|-----|
| Shapes: 35682  752  87  17  3  0  0  0 |
| % Total: 97%  2%  0%  0%  0%  0%  0%  0% |
|-----|

```

As the output shows, for each defined spatial index level, the following values and statistics are printed:

- Grid level and cell size.
- Total spatial index records for the current grid level.
- Total geometries stored for the current grid level.
- Ratio of spatial index records per geometry.
- Geometry counts and percentages by group that indicate how geometries are grouped within the spatial index at this grid level. The column headings have the following meaning (where ‘N’ is the number of grid cells):

<=N Number of geometries and percentage of total geometries that fall within <= N grid cells

>N Number of geometries and percentage of total geometries that fall within > N grid cells

Notice that the ‘>’ groupings include count values from the next group. For instance, the ‘>4’ group count represents the number of geometries that require more than four grid records as well as more than 10, and so on.

- Average number of geometries per grid.
- Maximum number of geometries per grid. This is the maximum number of geometries indexed into a single grid.
- Percentage of geometries wholly inside one grid. This is the percentage of all geometries wholly contained by one grid record.

The output sample shows spatial index statistics for a map feature class that uses two grid levels: one that specifies a grid size of 200 meters, the other a grid size of 1,600 meters. When a geometry requires more than four spatial index records, it is automatically promoted to the next grid level, if one is defined. That is, in no case will a geometry generate more than four grid records if more than one grid level exists. If a higher grid level doesn’t exist, then a geometry can have more than four grid records.


In the example above, 627,392 features are indexed through grid level 1. Because the system automatically promotes geometries that need more than four spatial index records to the next defined grid level, all 627,392 geometries for grid level 1 are indexed with four grid records or less. Grid level 2 is the last defined grid level, so geometries indexed at this level are allowed to be indexed with more than four grid records. At grid level 2,

there are a total of 36,434 geometries and 70,532 spatial index records; 35,682 geometries are indexed with four grid records or fewer, 752 geometries with more than four grid records, 87 geometries with more than 10, 17 geometries with more than 25, and three geometries with more than 50 grid records. Percentage values below each column show how the geometries are dispersed through the eight groups.

Indexes

Every data table in ArcSDE for MS SQL Server has a clustered index. The clustered key in an ArcSDE table is almost always the join key in a query like the business table's spatial column or the *f* table's *fid* column. The SQL Server query processor favors clustered indices because a clustered index orders the table data at its leaf level based on the index key. A normal index would have a pointer to the data page. Clustered indexes are b-tree structures containing pages with index rows holding a key value and a pointer to either a data page or a lower-level index page. All searches begin at the top level of the b-tree, the root node. This page is found in the *sysindexes* table in the *root* field. Within this page are pointers to intermediate pages, themselves holding ranges of values and pointers to other pages.

Figure A.2: B-tree diagram of a clustered index



Here's an example of the contents of the root node (shortened for brevity's sake). The Page Header displays 0:0 for both *m_nextPage* and *m_prevPage*, indicating that it is a root node. Also, two data slots are displayed, each with a key value and a pointer to another page. The pages pointed to are intermediate pages, themselves containing more key values and pointers.

PAGE HEADER:

Page @0x19A8A000

m_pageId = (1:4713)
m_typeFlagBits = 0x0
m_objId = 1526296497
m_nextPage = (0:0)
m_freeCnt = 5968
m_lsn = (69:1958:6)
m_ghostRecCnt = 0

m_headerVersion = 1
m_level = 0
m_indexId = 1
pminlen = 11
m_freeData = 1958
m_xactReserved = 0
m_tornBits = 0

m_type = 2
m_flagBits = 0x0
m_prevPage = (0:0)
m_slotCnt = 133
m_reservedCnt = 0
m_xdesId = (0:0)

DATA:

```

Slot 0, Offset 0x60
-----
Record Type = INDEX_RECORD
Record Attributes = NULL_BITMAP
19a8a060: 00000116 00126800 02000100 0000 .....h.....

Slot 1, Offset 0x6e
-----
Record Type = INDEX_RECORD
Record Attributes = NULL_BITMAP
19a8a06e: 00001816 00126a00 02000100 0000 .....j.....

```

Since a cluster orders the storage of a table, you cannot specify different filegroups for the cluster and its table. For example, if you specify in the dbtune table that the feature table's clustered index should go on filegroup `f_idx` and the table should reside on `f_storage`, your table will be stored with the index on `f_idx`.

What will happen here?

DEFAULTS	F_INDEX_CLUSTER	1
DEFAULTS	F_INDEX_FID	WITH FILLFACTOR=75 ON F_IDX
DEFAULTS	F_STORAGE	ON F_STORAGE

Index maintenance

When you initially create a table with a clustered index and specify a fill factor for that index, each data page is filled to that fill percentage when you initially insert rows. Later, when you edit this table with insert, delete, and update statements, the amount of data contained in each page will tend to increase.

If a data page fills to 100 percent and that particular page incurs an update or an insert, the page splits approximately in half. Half of the page moves into a newly allocated extent not contiguous to its former location. This will fragment your tables and increase query time as the disk spindles must travel to more locations to fetch data.

You can use the fillfactor clause in the dbtune table to delay page splits by limiting the fill percentage of a data page when it is initially allocated. Any ensuing inserts or updates will add to the fill percentage of a data page. Monitor page splits with DBCC SHOWCONTIG, and when your tables become fragmented, rebuild your clustered indexes or run DBCC INDEXDEFRAG (SQL Server 2000 only) or DBCC DBREINDEX.

Rebuilding a clustered index will reorder its table and defragment it. Data pages are reset to their initial fillfactor setting. The simplest way to rebuild an index is to place an ArcSDE feature class into Load_only_io and then back to normal_io. However, this won't work on versioned data.

If your data is multiversioned, you don't need to worry about monitoring the ArcSDE business table for extent fragmentation. Instead, your delta tables (a&d), feature tables, and spatial index will become fragmented. You can issue a DBCC DBREINDEX statement or invoke the DBCC INDEXDEFRAG command.

Here's an example of DBCC SHOWCONTIG run against a feature class that has incurred several edits—in this case, many updates and deletes. The output shown below reflects only fragmentation against the spatial index.

```
dbcc showcontig('hydro.s5')

DBCC SHOWCONTIG scanning 's5' table...
Table: 's5' (1622296839); index ID: 1, database ID: 7
TABLE level scan performed.
- Pages Scanned.....: 57
- Extents Scanned.....: 11
- Extent Switches.....: 36
- Avg. Pages per Extent.....: 5.2
- Scan Density [Best Count:Actual Count].....: 21.62% [8:37]
- Logical Scan Fragmentation .....: 33.33%
- Extent Scan Fragmentation .....: 27.27%
- Avg. Bytes Free per Page.....: 3780.6
- Avg. Page Density (full).....: 53.29%
```

For an exact explanation of these numbers, see the SQL Server Books Online under DBCC SHOWCONTIG. In this case, 36 extent switches are far too many to cover 57 pages. The best count reflects an optimal amount of extent switches, which in this case would be 8. Both extent scan fragmentation and logical scan fragmentation should be as close to 0 as possible. These measurements reflect the ordering of pages and extents in the index allocation map. Running either of these commands on this table will defragment the table.

```
dbcc indexdefrag (sde,'vtest.s5',s5_pk)
dbcc dbreindex('hydro.s5')

DBCC SHOWCONTIG scanning 's5' table...
Table: 's5' (1622296839); index ID: 1, database ID: 7
TABLE level scan performed.
- Pages Scanned.....: 41
- Extents Scanned.....: 6
- Extent Switches.....: 5
- Avg. Pages per Extent.....: 6.8
- Scan Density [Best Count:Actual Count].....: 100.00% [6:6]
- Logical Scan Fragmentation .....: 4.88%
- Extent Scan Fragmentation .....: 33.33%
- Avg. Bytes Free per Page.....: 2096.6
- Avg. Page Density (full).....: 74.10%
```

DBCC DBREINDEX and DBCC INDEXDEFRAG differ in that INDEXDEFRAG can be run online. It will not block queries or updates. For relatively unfragmented tables, DBCC INDEXDEFRAG should run faster than DBREINDEX. For more severely fragmented data, use DBREINDEX.

You only need to worry about business table fragmentation after running a full compress, when the delta tables have been flushed into the business table. Therefore, you should rebuild your business tables' clustered indexes after running a full compress.

NETWORK PACKET SIZE (SDEPACKETSIZE)

ArcSDE for MS SQL Server stores geometry in an image data type column, points, of the feature table. Several Geodatabase network tables use image type columns. Microsoft recommends increasing the size of the **network packet size** setting when employing image data type columns.

Network Packet Size is the size of the tabular data scheme (TDS) packets used to communicate between applications and the relational database engine. The default packet size is 4 KB and is controlled by the network packet size configuration option.

ArcSDE, by default, sets this to 8192k, double its default setting of 4096k. You can also make this setting global to your MS SQL Server by using the sp_configure statement's "network packet size" setting.

```
sp_configure 'show advanced options',1
```

```

reconfigure with override

go

sp_configure 'network packet size',8192

reconfigure with override

go

```

You can control the network packet size setting with the environment variable SDEPACKETSIZE, although we recommend you stay with the default setting.

Text in row (SQL Server 2000 Only)

SQL Server 2000 allows up to 7 KB of image, text, or ntext data to be stored directly in row on a data page. SQL Server 7 stored these data types in separate image pages, located with a 16-byte pointer in the table's data page. Text in row at SQL Server 2000 permits you to place image type data directly in the data page, and anything exceeding the setting is placed in a traditional image page. All ArcSDE geometry is stored in an image type column in the feature table. ArcSDE Raster layers and Geodatabase network tables use image data types.

This setting is enabled on tables using the command `sp_tableoption`. ArcSDE does this automatically for you through the `dbtune` table. The `SDE_dbtune` table sets a default of 256 bytes for many settings.

Keyword	Parameter	Explanation
DEFAULTS	B_TEXT_IN_ROW	Business tables with BLOB data type
DEFAULTS	F_TEXT_IN_ROW	Feature tables points column
NETWORK_DEFAULTS	B_TEXT_IN_ROW	Network business table with BLOB type
NETWORK_DEFAULTS	F_TEXT_IN_ROW	Feature table for generated junctions class
NETWORK_DEFAULTS::NETWORK	B_TEXT_IN_ROW	N_ tables with image type columns

This may not be the best size for all your data. You don't want to allow too much text in row space on your data pages if you can't fill it. For example, it would be unwise to set this option to 7000 because you would create many more data pages and extents, thus lengthening the distance a query must travel to retrieve data. Likewise, underestimating this number would create more image data pages and could produce the same problem.

One way to determine an optimal table text in row setting is to estimate the 80th% size of all your image type records and store up to that size in the `*_text_in_row` setting. First you'll have to load your data and then run this query:

```

select top 20 percent (datalength(points)) from boris.f12 order by
datalength(points) desc

```

This query will list the top 20 percent largest data records in the "points" column of a feature class's feature table. You then pick the smallest value from its result set and use

that for a dbtune entry and reload your data, referencing that keyword. In this case, the 80th% value equaled 455, so you would make this entry in the dbtune:

Keyword	Parameter	Explanation
CARTO	F_INDEX_FID	WITH FILLFACTOR=75 ON FG_CARTO
CARTO	F_TEXT_IN_ROW	455

Now reload this data referencing this keyword to pick up the new F_TEXT_IN_ROW setting. You could also put this value in the DEFAULTS keyword, but then all data loads would use this size.

```
Shp2sde -o create -l silos,shape -f silos -a all -D carto -s bigboy -u
boris -p badinoff -k CARTO
```

You may want to check if a particular table has been enabled for text in row and what its current size is set to. Use the OBJECTPROPERTY Transact-SQL metadata function:

```
select objectproperty(object_id('boris.fl2'),'tabletextinrowlimit')
```

You may not want to reload your data. In this case, you can enable a table for text in row with sp_tableoption and write a query to copy all the existing image data types into a new column:

```
/*
FeatureTable_Text_in_Row.sql
script to copy existing points data into new column
with new text in row value enabled

SQL Server 2000 Only

Script Unsupported by ESRI
*/

sp_tableoption 'boris.fl2','text in row',455
go
alter table boris.fl2
add newpoints image
go
update boris.fl2
set newpoints=points
go
alter table boris.fl2
drop column points
go
sp_rename 'boris.fl2.newpoints','points','column'
go
```

If you have a lot of existing data in SQL Server 7.0 and desire to move this data into SQL Server 2000, you'll want to move this data into the new text_in_row format. How you perform your upgrade will determine how your data is moved from a traditional image page into text_in_row.

Text_in_row is enabled with the sp_tableoption command, and ArcSDE sets a text_in_row default (in the sde_dbtune table) for all feature tables to 255 bytes. This means that any layer or feature class creation in ArcSDE on SQL Server 2000 will use the text_in_row option. To upgrade an existing SQL Server 7.0 database to SQL Server 2000 and push your image data into text_in_row format, you can:

- € Perform a Geodatabase copy and paste between two ArcSDE servers: one running SQL Server 7.0 and the other SQL Server 2000.
- € Reload all your data into ArcSDE on SQL Server 2000.

- € Run the SQL Server 2000 upgrade wizard or restore your SQL Server 7.0 databases to a SQL Server 2000 instance. Once you have done this, run the `FeatureTable_Text_in_Row.sql` listed above.

TEXT_IN_ROW for rasters

A raster stored in ArcSDE is composed of four tables. Of those, the `sde_blk_<n>` table holds the largest image type column. You can enable `text_in_row` for this table with the dbtune entry `BLK_TEXT_IN_ROW`. For example:

RASTER	BLK_TEXT_IN_ROW	2000
--------	-----------------	------

We recommend you stay with the defaults as setting a large `text_in_row` will create too many data pages and extents, eventually slowing your queries.

Triggers

All feature class business tables have a delete update trigger and an insert trigger. The insert trigger ensures that the spatial column cannot have duplicate values in the business table, while the delete-update trigger manages activity against the spatial column in the business, feature, and spatial index tables. You can view these triggers in the SQL Server Enterprise Manager.

These triggers are automatically dropped whenever a feature class, either standalone or in a feature dataset, is multiversioned. They are re-created when a feature class is “unregistered as versioned.” If you edit one of these triggers and then multiversion the feature class, the trigger will be dropped.

“I” tables and stored procedures

You’ll notice that once you’ve loaded data or run `sdesetupmssql`, you will have several “i” tables and stored procedures in your databases. These stored procedures and tables are used for fetching feature IDs for feature classes. Editing these tables or stored procedures is not supported and highly discouraged.

APPENDIX B

The components of the installation

This *ArcSDE Configuration and Tuning Guide for Microsoft SQL Server* covers material after the installation of the ArcSDE software. Users should now be prepared to start the ArcSDE service or make direct client connections to the spatial data store. If the ArcSDE software has not been installed, refer to the *ArcSDE Installation Manual* for directions. This chapter describes what the installation configured on your hard disk, the services it created, what must be present for a successful ArcSDE service start, and how to start the ArcSDE service. Direct connection two-tier ArcSDE setup will not be covered in this chapter; this is covered in Chapter 5, ‘Connecting to SQL Server.’

ArcSDE installation basics

The ArcSDE installation moves several files to your hard disk, sets one environment variable, modifies the PATH environment variable, edits two system files, and creates several registry entries. To successfully start the ArcSDE services, there must be:

1. SDE Service file edits:

The services.sde file should contain the name and TCP/IP port number of the ArcSDE service. This file is located in %sdehome%\etc. If the ArcSDE service is called “esri_sde” and resides on port 5151, then this file should read:

```
Esri_sde 5151/tcp
```

2. The Windows 2000 or Windows NT Services file:

This file, “services”, is located in %windir%\system32\drivers\etc and must match the entry in point 1. There must be a carriage return after the last line in this file and its counterpart in %sdehome%\etc.

3. Windows Registry:

The ArcSDE Instance should be listed in the windows registry under these keys:

HKEY_LOCAL_MACHINE\Software\ESRI\ArcInfo\ArcSDE\8.0\ArcSDE for SQL Server

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services

HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Services

HKEY_LOCAL_MACHINE\SYSTEM\ControlSet002\Services

The entries in these keys are vitally important. They must be correct for a successful ArcSDE service start. These entries are discussed below in the section entitled “Manually deleting and creating an ArcSDE service.”

4. The environment variable %SDEHOME% must point to the installation directory of the ArcSDE software. There can only be one path, and it must contain the ArcSDE bin, etc, lib, tools, and codepage folders.

%SDEHOME%=d:\ArcSDE81

5. There must be an ‘sde’ database and an ‘sde’ user. The sde user can either be SQL Server or Windows NT/Windows 2000 authenticated. Either way, this user must have at least create table and create procedure within the ‘sde’ database.

6. There must be a valid license. Check in the sdelic.log file in %sdehome%\etc. If the service is properly licensed, there will be an entry similar to:

```
Initializing connection to license server (@ZAPHOD)
```

7. The sde and gdb metadata tables must be present in the database, and they must be owned by the sde user. There should be at least 20 ‘SDE_’ tables (not counting the sde_logfile* tables) and 26 ‘GDB_’ tables present. There should also be at least 66 ‘SDE_’ stored procedures in the sde database. The install creates the metadata tables by running the sdesetupmssql.exe executable. See Chapter 5, ‘Connecting to SQL Server’, for instructions on running this command.

If the seven points above are correct and the ArcSDE service will not start, add the line SET SDEVERBOSE=TRUE in the dbinit.sde file in %sdehome%\etc and try again to start the ArcSDE service. If it fails to start, open the sde.errlog in %sdehome%\etc. There will appear an informative written error message here.

How to start the ArcSDE service

After verification that the above procedure has been correctly performed, it is now possible to start the ArcSDE service to accept incoming connection requests from client programs. If a direct connection to ArcSDE without using the three-tier architecture is desired, refer to Chapter 5, 'Connecting to SQL Server.'

To start the ArcSDE service, there must be a Windows NT or Windows 2000 named 'service.' Use one of the three tools described below to start this service.

Use the "services" utility in the control panel or services ui. Open the services form in the control panel, locate the ArcSDE service name configured during installation, select it, and press the Start button.

Use the "net" command from a command prompt. Open a command prompt and type: net start <service name>

```
c:\ArcSDE81\bin\net start esri_sde
```

Use the sdemon tool located in %SDEHOME%\bin. Open a command prompt and type: sdemon -o start -i <service name> -p <sde dba pass>

```
c:\ArcSDE81\bin\sdemon -o start -i esri_sde -p big.bob
```

How to stop the ArcSDE service

Once the ArcSDE service is successfully started, it can be stopped by employing one of the following methods:

Use the services utility in either the Windows NT control panel or the Windows 2000 services management console under System Administration. Locate the ArcSDE service, select it, and click the stop control.

Use the "net" command from a command prompt. Open a command prompt and type "net stop <service_name>".

```
c:\ArcSDE81\bin\net stop esri_sde
```

Use the sdemon command located in %SDEHOME%\bin. Open a command prompt and type: sdemon -o shutdown -i <service name> -p <sde admin password>

```
c:\ArcSDE81\bin\sdaemon -o shutdown -i esri_sde -p
big.bob
```

Manually deleting and creating an ArcSDE service

The ArcSDE installation creates an ArcSDE service. The installation prompted the user for a service name, TCP/IP port number, license manager, and the sde dba password.

However, once the ArcSDE binaries have been installed on disk, it is possible to create a new service manually using the 'sdeservice' tool in %SDEHOME%\bin.

Multiple ArcSDE services are required if there are multiple instances of Microsoft SQL Server running simultaneously on a single server. Microsoft SQL Server 2000 allows multiple installations of the database software on a single server. Similarly, there can be one MS SQL Server 6.5 instance or 7.0 instance running in conjunction with the MS SQL Server 2000 instance. This can be valuable in upgrading a SQL Server 7.0 instance to SQL Server 2000. This tool can also be run as a substitute for uninstalling/reinstalling the ArcSDE software. If it is believed there is a corrupted instance, delete the ArcSDE service by running sdeservice -o delete, then re-create it with sdeservice -o create.

To remove an ArcSDE service:

```
sdeservice -o delete -i esri_sde -p big.bob
```

Running the above command will remove the service entry from the Windows Registry in these keys:

- ⊘ **\\HKEY_LOCAL_MACHINE\SOFTWARE\ESRI\ArcInfo\ArcSDE\8.0\ArcSDE for SQLServer**
- ⊘ **\\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\<service name>**
- ⊘ **\\HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001**
- ⊘ **\\HKEY_LOCAL_MACHINE\SYSTEM\ControlSet002**

The minimum required arguments to create a new service are service name (-i esri_sde), license server name (-l zaphod), and sde dba password (-p big.bob).

```
Sdeservice -o create -p big.bob -l zaphod -i esri_sde
```

The above command will re-create all the pertinent registry keys. NOTE: This command's usage has a [-u <service_user>] [-P <service_user_password>] switch. For example:

```
sdeservice -o create -p <DB_Admin_password>  
-l <license server name>  
[-q] [-H <sde_directory>]  
[-d  
<ORACLE,SID|ORACLE8I,SID|SQLSERVER|DB2,DB2INSTANCE|INFORMIX>]  
[-i <service>]  
[-u <service_user>] [-P <service_user_password>]  
[-s <OLE DB datasource (default: local hostname)>]
```

Use of the -u and -P arguments is not necessary. Use of these switches specifies a domain user and password, not a SQL Server user or password. The ArcSDE install program does not set these switches. Care should be exercised in the use of these switches. The -p argument specifies the sde user's password.

The “DATASOURCE” of an ArcSDE service

The sdeservice command has a DATASOURCE argument with the -s switch. The command sdesetupmssql -o upgrade also has a DATASOURCE argument. A datasource represents a named SQL Server instance. At SQL Server 7.0, the datasource is generally the host name of the server hosting the database.

For example, to create an ArcSDE service specifying a datasource, use the -s switch followed by the name of the database server host. In this case, the datasource is “vogon”, meaning the name of the server hosting SQL Server 7 is “vogon.” The license manager resides on the server “zaphod.”

```
Sdeservice -o create -p big.bob -l zaphod -i esri_sde -s  
vogon
```

Using DATASOURCE to separate ArcSDE and the SQL Server host

Although this would probably result in detrimental performance, it is possible to install the ArcSDE software onto a server other than the SQL Server. During the install or when invoking the sdeservice command, simply specify the datasource in the -s switch. For example, it is desired to point to the ArcSDE datasource for testing

purposes at another SQL Server. The ArcSDE application software is located on the server “vogon”, but its datasource points to another SQL Server, “marvin.”

```
sdeservice -o create -p big.bob -l zaphod -i esri_test -s marvin
```

When starting this new service, the data store holding ArcSDE tables will actually be located on another machine, marvin.

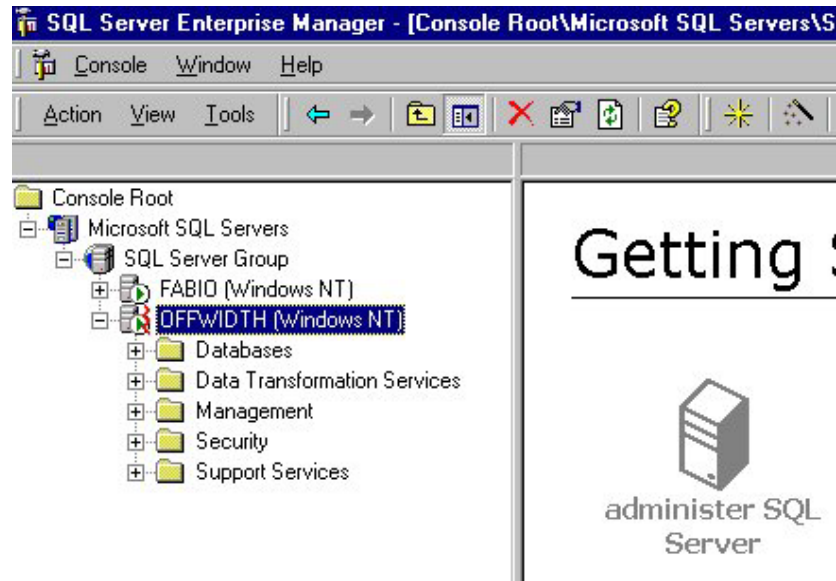
Using DATASOURCE with SQL Server 2000

During the installation of SQL Server 2000, the user is prompted to enter an instance name. This instance name becomes the datasource parameter for ArcSDE. Also, to run SQL Server 7 and SQL Server 2000 on the same machine, it is necessary to provide different instance names so that two different ArcSDE instances can be created for them. For example, if SQL Server 2000 is installed on server “Vogon,” the instance name on Vogon for SQL Server 2000 becomes “Vogon\SQL2k.” To create two ArcSDE services on Vogon for SQL Server 7 and SQL Server 2000, you would use:

```
sdeservice -o create -i esri_sql -p big.bob -l zaphod -s vogon /*SQL Server 7*/  
sdeservice -o create -i esri_sql -p bob_2k -l zaphod -s vogon\sql2k /*SQL Server 2000*/
```

Their datasource will be listed in the SQL Server Enterprise Manager under the “SQL Server Group” branch. In the example below, the highlighted name “TETHYS” would be the datasource name for the server.

Figure B.1: The highlighted entry “OFFWIDTH” represents a datasource.



To create an instance with TETHYS as the datasource, the service creation command would look like this:

```
sdeservice -o create -i esri_sql -p big.bob -l
theodolite -s Tethys
```

Upgrading SDE 3.0.2 or ArcSDE 8.0.2 to ArcSDE 8.3

If either SDE 3.0.2 or ArcSDE 8.0.2 is installed, the ArcSDE 8.3 installation will attempt to upgrade it to ArcSDE 8.3. **NOTE:** An upgrade requires the use of a sysadmin account.

When the install detects an ArcSDE 8.0.2 instance on the server, it makes an internal connection to SQL Server as a sysadmin fixed server role user, creates new stored procedures for each ArcSDE table in the database, then grants execute on those stored procedures to all users who have permissions to the data. New tables at ArcSDE 8.3 are also created. All data owners are granted create procedure as well. The upgrade process changes the sde user's permissions in the sde database to create table, view, stored procedure, and function (SQL Server 2000 only).

When the install detects an SDE 3.0.2 or 3.0.2.1 instance, it restructures the SDE “F” tables by moving all data stored in the varbinary columns into the points image column. Then it upgrades all layers and metadata to the ArcSDE 8.3 format. Once

the layer migration is complete, ArcSDE rebuilds all indexes on its data. Upgrading from SDE 3.0.2 to ArcSDE 8.3 will probably increase the size of your database.

It is also possible to manually upgrade an 8.0.2 or 3.0.2/3.0.2.1 database to ArcSDE 8.3 with the `sdesetupmssql.exe` command in `%SDEHOME%\bin`. The upgrade option requires the command to be run as a sysadmin user.

To upgrade an ArcSDE 8.0.2 instance to ArcSDE 8.3

Install ArcSDE 8.3. The installer will detect the existence of ArcSDE 8.0.2 and attempt to upgrade an 8.0.2 schema.

To upgrade, the installer runs the command “`sdesetupmssql -o upgrade`.” This command can also be run manually. The command must be run by a sysadmin account.

After the command finishes, check the `upgrade.log` file in `%SDEHOME%\etc`. If the file does not end with a report of success, rerun the command manually until it does. **NOTE:** If you have a customized `dbtune.sde` file from ArcSDE 8.0.2, this file must be resident in the `%SDEHOME%\etc` folder to successfully run `sdesetupmssql -o upgrade`. The installer does present you with the option of copying an existing `dbtune.sde` into the new installation.

To upgrade an SDE 3.0.2 or SDE 3.0.2.1 instance to ArcSDE 8.3

This example presumes that you will upgrade an SDE 3.x instance running on SQL Server 7.0 to ArcSDE 8.3.

Shutdown your SDE 3.x instance.

Install MDAC 2.6 or greater. Microsoft Data Access Components (MDAC) 2.6 is available through Microsoft.

Install a supported SQL Server Service Pack. Service packs are available through Microsoft.

Install ArcSDE 8.3. The installer will attempt to remove the version 3.x software. The postinstallation will attempt to upgrade the version 3.x instance to version 8.3. The upgrade must be run as a sysadmin account.

The upgrade will take some time and depends upon how much data you have and how large your tables are. If the upgrade returns an error, check the `upgrade.log` file in `%SDEHOME%\etc`, correct the file, and start the command (`sdesetupmssql -o upgrade`) again. This command is located in `%SDEHOME%\bin`. The `-o upgrade` option must be run as a sysadmin account. Rerun the command manually until it does. The installer does present you with the option of copying an existing `dbtune.sde` into the new installation.

Upgrading an ArcSDE instance on SQL Server 7 to SQL Server 2000

It is recommended that users first allow the installation to upgrade any existing SDE 3.0.2.x or ArcSDE 8.0.2 instances to ArcSDE 8.3 before migrating data into SQL Server 2000. ArcSDE 8.3 leverages the SQL Server 2000 Text in Row option for each table with an image data type. Since all F tables and many network tables use image types, it is beneficial to use this setting. See the SQL Server 2000 Books Online for

information regarding `sp_tableoption`'s 'text in row' setting. The option, as it relates to ArcSDE 8.3, is discussed in Appendix A.

To upgrade ArcSDE from SQL Server 7 to SQL Server 2000

There are two basic options:

Allow the ArcSDE install to upgrade the ArcSDE 8.0.2 or SDE 3.x instance to version 8.3 or do this manually with the `sdesetupmssql` tool in `%sdehome%\bin`. To perform an upgrade, the user must initiate this command as a member of the `sysadmin` fixed server role. This can be either the 'sa' user or an Nt/Win2k user.

If you have the licensing available, create an instance of ArcSDE 8.3 on SQL Server 2000. Next, if ArcGIS is installed, use the Geodatabase copy and paste tool to copy the data from the version 8.3 instance on SQL Server 7 to the version 8.3 instance on SQL Server 2000. Using Geodatabase copy and paste will preserve relationship classes, geometric networks, and other advanced Geodatabase models. This will require starting two instances of ArcSDE. The image data types (all ArcSDE geometry) will be created with a default text in row setting of 255. The new `dbtune` table sets this automatically.

If ArcGIS is not installed or you don't have the licensing to maintain two instances of ArcSDE, use SQL Server tools such as backup and restore, `sp_detach_db` and `sp_attach_db`, DTS, or the upgrade wizard. However, the image data will not be upgraded to the text in row format. Appendix A shows a t-sql example of moving image type data into the SQL Server 2000 `text_in_row` format.

When upgrading from SQL Server 7 to SQL Server 2000 use the `sp_dbcmptlevel` stored procedure to set certain database behaviors to be compatible with the specified earlier version of SQL Server. A server upgraded from SQL Server 7 to SQL Server 2000 will have all its databases compatibility level set to level 80. If a server is upgraded from SQL Server 6.0 or 6.5 will have its databases set to the SQL Server version which the database was last used, either 60 or 65. All databases in SQL Server 2000 must have their compatibility level set to 80. To verify that all databases are set to compatibility level 80, execute the following procedure:

```
EXEC sp_dbcmptlevel 'database name'
```

To change to a different compatibility mode, run the following command:

```
EXEC sp_dbcmptlevel 'database name' compatibility-mode
```

The `sp_dbcmptlevel` command is covered in greater detail in SQL Server Books Online.

In addition the `sdesetupmssql -o upgrade` command must be run on upgraded instances of SQL Server 7 to SQL Server 2000. To perform an upgrade, the user must initiate this command as a member of the `sysadmin` fixed server role. This can be either the 'sa' user or an Nt/Win2k user. The syntax of this command is:

```
-o upgrade [-H <sde_directory>] [-u <DB_Admin_user>] [-p <DB_Admin_password>]  
          [-s <datasource>] [-N] [-q]
```

Upgrade examples

NOTE: These are simple examples of what you can do to upgrade your instance. There are many ways to do this, the most simple of which, and the only supported method, is to let the installer upgrade an 8.0.2 instance to 8.3 or manually run the `sdesetupmssql -o upgrade` command.

Example one: ArcSDE 8.0.2 on MS SQL Server 7.0 to ArcSDE 8.3 on SQL Server 2000 (different servers).

Install SQL Server 2000 and ArcSDE 8.3 onto the new server.

Create the same databases and users as presently configured on the SQL Server 7.0, ArcSDE 8.0.2 instance.

Start an ArcSDE 8.3 service on the SQL Server 2000 server.

Use the copy and paste tool from ArcCatalog to move the data from the 8.0.2 instance on SQL Server 7.0 to the ArcSDE 8.3 instance on SQL Server 2000.

Note: You must have enough ArcSDEServer licenses to do this.

After the copy and paste has finished, backup the database.

Example two: ArcSDE 8.0.2 on MS SQL Server 7.0 to ArcSDE 8.3 on SQL Server 2000 (same server)

Install ArcSDE 8.3 onto the server. The installer upgrades the version 8.0.2 instance to ArcSDE 8.3. The installer will execute the command “sdesetupmssql –o upgrade” as a sysadmin user. You must supply this username and password during the postinstallation portion of the install. This command is located in the %SDEHOME%\bin folder of the ArcSDE 8.3 installation.

Start the ArcSDE 8.3 service to verify that the upgrade succeeded. If the service will not start, check the “upgrade.log” file in %SDEHOME%\etc for information. Correct any errors listed in this log and manually re-execute the sdesetupmssql –o upgrade command as a sysadmin user. Repeat this process until the upgrade succeeds.

Install SQL Server 2000 onto the server. Shutdown the ArcSDE 8.0.2 service. Use a SQL Server tool to move the 7.0 databases to the 2000 server. You can use backup and restore, sp_detach_db and sp_attach_db, DTS, or the SQL Server 2000 Copy Database Wizard to do this. See the SQL Server 2000 Books Online for this information and Chapter 8, ‘Backup and recovery’, of this guide for an example of using backup and restore to move databases. If you restore a backup of your SQL Server 7.0 databases to a SQL Server 2000 instance, you must manually execute the sp_dbcmptlevel stored procedure to set your database’s compatibility level to ‘80’ (SQL Server 2000). You must do this for all restored databases. For example:

```
sp_dbcmptlevel sde, '80'
```

Modify your existing ArcSDE 8.3 service to point to the SQL Server 2000 datasource. Use the sdeservice –o modify tool to change the –r DATASOURCE argument of your service. You can also delete and re-create your service. See the section above, Manually deleting and creating an ArcSDE service, for information on how to do this.

Start your ArcSDE 8.3 service.

Example three: ArcSDE 8.0.2 on MS SQL Server 7.0 to ArcSDE 8.3 on SQL Server 2000 (same server, maintaining both the version 8.0.2 instance and version 8.3 instance). This upgrade scenario is more involved. In order to maintain both the older service and the version 8.3 service on the same machine, you’ll need to do the following:

Install SQL Server 2000 onto the server. Installing SQL Server 2000 onto the same machine as SQL Server 7.0 is possible with SQL Server 2000’s support of multiple instances. See the section in this guide on multi-instance support for more information.

Before installing ArcSDE 8.3, shutdown ArcSDE 8.0.2. Make a copy of your older SDEHOME (all the directories and files located beneath SDEHOME). You will need to copy it back after the 8.3 installation is finished.

Install ArcSDE 8.3. It will remove your ArcSDE 8.0.2 installation, delete your ArcSDE 8.0.2 service, and remove your port entry in the %windir%\system32\drivers\etc\services file. Do not drop the sde database. SDEHOME will be set to the 8.3 install directory, and your path will include the 8.3 install directory\bin folder.

Start the ArcSDE 8.3 postinstallation. **Click Next>** in the *Welcome Screen*.

Click Next> in the *Define User Environment Screen* to create the sde login, and sde database and to add the login to the database as a user. If you precreated the sde login and database and have granted the login access to the sde database, **press Skip** in this screen and **click Next** in *Setup Configuration Files Screen* and go to step 8.

In the next form, input a sysadmin user and password and your SQL Server 2000 datasource. The wizard will proceed through the steps of creating an sde login and database. The following screenshot displays a SQL Server 2000 datasource on a machine called "Rift."

Click Next> through the wizard until you get to the *Repository Setup* form.

The *Repository Setup* form will create the sde metadata if it does not exist or upgrade a version 3.0.2.x sde or version 8.0.2 ArcSDE instance if it is present. **Click Next>**. **Click Skip** if your version 8.3 metadata already exists and proceed to step 10.

Now connect as a sysadmin user and input your SQL Server 2000 datasource. Again, see the above screenshot for how this should look. **NOTE: If you input your SQL Server 7.0 Datasource here, you will upgrade your version 8.0.2 instance to version 8.3** Finish the version 8.3 Postinstallation process of creating and starting a service.

Open a DOS prompt and change directories into the ArcSDE 8.0.2 %SDEHOME%\bin folder. Re-create your ArcSDE 8.0.2 service with the sdeservice -o create command. Make sure you get the datasource argument correct; your datasource should be the name of the SQL Server 7.0 server. **NOTE: You cannot point to two instances of ArcSDE at the same datasource.** Don't forget to use the -H flag of the sdeservice -o create command to register the correct SDEHOME for version 8.0.2.

```
C:\ArcSDE802\bin\sdeservice -o create -p <sde user password in ss7> -l
<license server> -H c:\ArcSDE802 -i esri_sql802 -s <server name>
```

Open the services file in %WINDIR%\system32\drivers\etc and readd your ArcSDE 802 port. For example, you'll need to readd this line:

```
esri_sql802          5151/tcp
```

Before starting both services, make sure your services are correct by running `sdeservice -o list` within the bin folder of each SDEHOME, for example, for ArcSDE 8.0.2:

```
C:\ArcSDE802\bin>sdeservice -o list
```

```
SDE service Information
```

```
-----
Service Name:          ArcSde Service(esri_sql802)
Service SDEHOME:       C:\ArcSDE802
Service License Server: @theodolite
Service Datasource:    rift
Service Status:        SERVICE_STOPPED
```

For ArcSDE 8.3:

```
C:\sde81\bin>sdeservice -o list
```

```
SDE service Information
```

```
-----
Service Name:          ArcSde Service(sde81)
Service SDEHOME:       C:\sde81
Service License Server: @theodolite
Service Datasource:    rift\sql2k
Service Status:        SERVICE_STOPPED
```

Note that in this example, each service has its own SDEHOME. Also note that each service has its own datasource argument. The ArcSDE 8.0.2 service points to the **rft** datasource, while the ArcSDE 8.3 service points to the **rft\sql2k** SQL Server 2000 datasource.

Now start one or both instances. **NOTE:** You'll have to have the proper licensing to do this. You can run both instances to test your data under the 8.3 configuration. You can move data from the 8.0.2 instance by copying and pasting with the ArcCatalog or exporting with `sdeexport`, `sde2shp`, etc.

Example four: ArcSDE 8.0.2 or SDE 3.0.2.x on SQL Server 7.0 to ArcSDE 8.3 on SQL Server 2000 (same machine)

Make backups of all your sde and spatial databases using the SQL Server 7.0 backup tool.

Shutdown your ArcSDE 8.0.2 or SDE 3.0.2.x instance. **Uninstall the SDE software. If you have customized dbtune.sde, you must save a copy of this file!**

Install SQL Server 2000. Create an identical set of logins on the SQL Server 2000 instance as exist on the SQL Server 7.0 instance.

Do one of the following three steps:

Restore all the SQL Server 7.0 databases to the SQL Server 2000 instance. Run `sp_dbcmptlevel` for each database you've restored:

```
sp_dbcmptlevel sde,'80'
```

or

Run `sp_change_users_login 'update_one','<user>','<user>'` for each user, in each, database. A detailed example of this is given in Chapter 8, 'Backup and recovery'.

Run `sp_detach_db` on the sde and any other spatial database in the SQL Server 7.0 instance. In the SQL Server 2000 instance, run `sp_attach_db` to reattach them to the newer server. Again, you must run `sp_change_users_login`.

or

Use the SQL Server 2000 Copy Database Wizard to copy the SQL Server 7.0 databases to the SQL Server 2000 instance.

Install ArcSDE 8.3. Run the postinstallation utility and let it upgrade your 8.0.2 databases from SQL Server 7.0 to ArcSDE 8.3 in SQL Server 2000. **NOTE:** If the upgrade utility throws an error, note the error, check the `upgrade.log` file in `%SDEHOME%\etc`, correct the listed error, and manually restart `sdesetupmssql -o upgrade`. `Sdesetupmssql -o upgrade` must be run as a sysadmin user. For example:

```
sdesetupmssql -o upgrade -u sa -p sa.password -s rift\sql2k -H
c:\sde81\sql.exe
```

You can run this command as many times as needed until the upgrade completes successfully. **If you customized your `dbtune.sde` file from 8.0.2, you must run the upgrade process, either from the postinstall or using the `sdesetupmssql -o upgrade` command, with the customized `dbtune.sde` file in `%sdehome%\etc`.** The installer will prompt you to use a different `dbtune.sde` file during the postinstallation.

7.0 Start your sde service: `Net start esri_sde`

Troubleshooting the ArcSDE service startup

Once you have the ArcSDE software on disk and Microsoft SQL Server installed, you should not have to remove either of them. Follow the process listed below, and you should be able to start your ArcSDE service once you're done.

1. Check for the existence of an 'sde' user and 'sde' database: Verify that the sde database exists and that the sde login is a user in this database by logging in as 'sde.' **Make note of the sde user's password.** The sde user must have at least create table, create view, and create procedure in the sde database. This can be verified by right-clicking the database in the Enterprise Manager and selecting Properties. Then click the Permissions tab and check what permissions the 'sde' user has in the sde database. Verify that the sde and gdb metadata tables exist by examining the sde database. There should be approximately 20 `SDE_*` tables, 26 `GDB_*` tables, and 66 `SDE_*` stored procedures. If these tables and stored procedures do not exist, they must be created by running the `sdesetupmssql.exe` command located in `%SDEHOME%\bin`. See Chapter 5, 'Connecting to SQL Server,' for instructions on running this command.

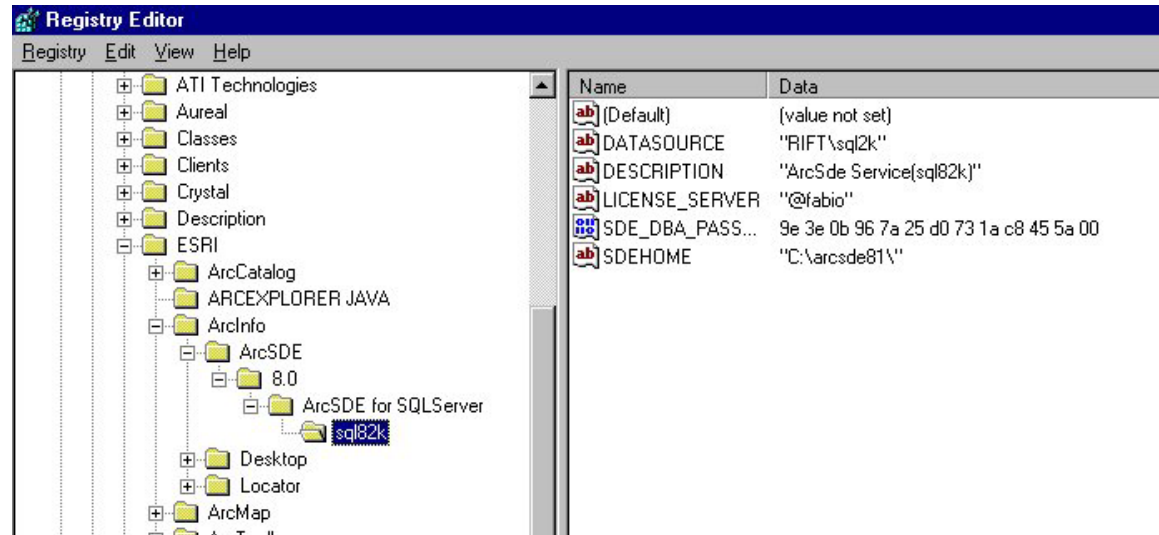
2. Check that the ArcSDE service is listed in both `%SDEHOME%\etc\services.sde` and `%WINDIR%\system32\drivers\etc\services`. The same service should be present in both. If this service is the last entry in these files, put a carriage return after it (blank space). The entries should look like the following:

```
Esri_sde 5151/tcp
```

3. The ArcSDE service must be correctly entered in the Windows Registry. The most critical locations in the registry are:

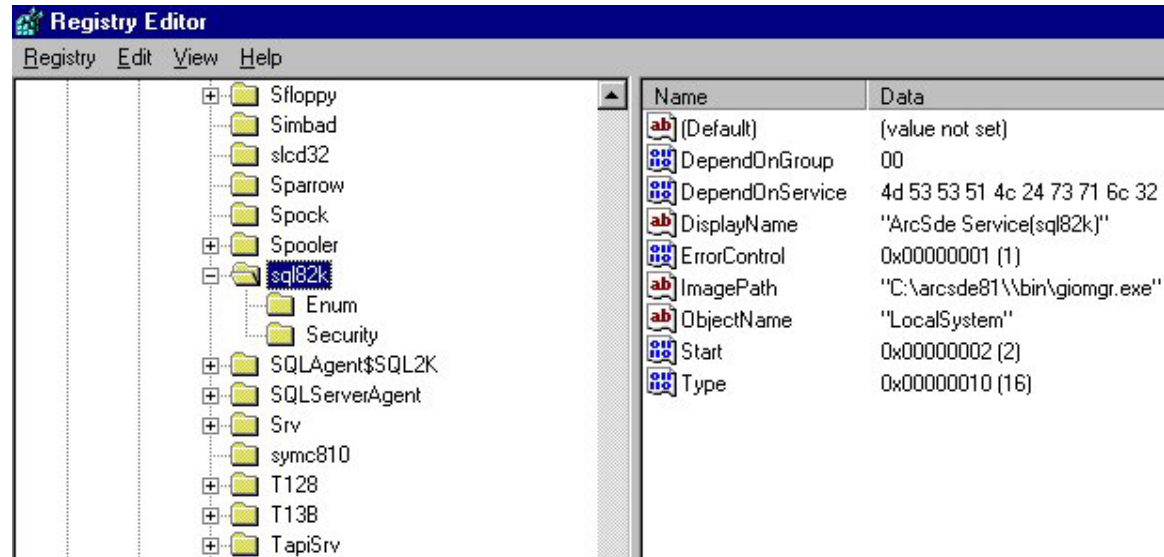
HKEY_LOCAL_MACHINE\Software\ESRI\ArcInfo\ArcSDE\8.0\ArcSDE for SQL Server—In this key, there will be the name of your service, its registered datasource, its license server, DBA_password, and SDEHOME. See Figure B.1. This service must be located under the “ArcSDE for SQL Server” key. Its DATASOURCE must point to your named SQL Server. A valid license_server should be listed with the @<name> convention. SDEHOME should point to an actual location on disk.

Figure B.2: A Service Entry in the HKEY_LOCAL_MACHINE\Software key



HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services—In this key, there will be a Windows representation of the actual service. See Figure B.2. The DependOnService represents the -d flag of the sdeservice -o create command and identifies a dependency between the ArcSDE service and another Windows NT service, most likely the MSSQLSERVER process. The ImagePath gives the path to the giomgr.exe executable, the io manager.

Figure B.3: A service entry in the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet key



Instead of editing the registry directly, delete the service with `sdeservice -o delete`. Since this is a DOS executable, be aware of the path from which it is run. Run it from `%SDEHOME%\bin` to ensure that it picks up the correct libraries. See the above section, 'Manually deleting and creating an ArcSDE service', for more information. Confirm that the password and data source arguments are entered correctly.

4. Confirm that there is a valid license manager process running on your network. The license manager may reside locally on the server, or it may be accessed on the network. You identify the name of the license host when installing or with the `-l` argument to the `sdeservice -o create` command. Verify your services' license server by typing `sdeservice -o list` from a DOS prompt:

```
C:\snoopyng>sdeservice -o list -i esri_sql
SDE service Information
```

```
-----
Service Name: ArcSde Service(esri_sql)
Service SDEHOME: C:\arcsde81\
Service License Server: @theodolite
Service Status: SERVICE_STOPPED
```

In this example, the service "esri_sql" has its license service running on theodolite. Theodolite is identified with the "@" symbol. Sometimes, you'll correctly setup the license server but for some reason not be able to access a license. In this case, ArcSDE appears to hang. You can tell if this has happened on Windows NT by examining the Windows NT Service Control panel's start and stop buttons. If both are grayed out, your service cannot access a valid license. You will also see an

‘unable to initialize connection to license server error’ message in the sdelic.log file in %SDEHOME%\etc.

5. In some rare instances, there will be an incorrect installation of MDAC, or it can become corrupted. MDAC is the Microsoft Data Access Components and is used to access the data stores in SQL Server through OLE DB. ArcSDE 8.3 supports MDAC 2.6 and greater. To verify which MDAC is installed, check the version of msado15.dll in program files\common files\system\ado or by downloading the component checker from Microsoft.

Now attempt to start the service. If it still cannot start, add the line “set SDEVERBOSE=TRUE” in the dbinit.sde file in %sdehome%\etc. The word ‘set’ must be in lower case. Now try again to start the service. Then open the sde.errlog file in %sdehome%\etc. There will be an informative error message written here.

APPENDIX C

Storing raster data

A raster is a rectangular array of equally spaced cells, which taken as a whole represent thematic, spectral, or picture data. Raster data can represent everything from qualities of land surface such as elevation or vegetation to satellite images, scanned maps, and photographs.

You are probably familiar with raster formats such as tagged image file format (TIFF), Joint Photographic Experts Group (JPEG), and Graphics Interchange Format (GIF) that your Internet browser renders. These rasters are composed of one or more bands. Each band is segmented into a grid of square pixels. Each pixel is assigned a value that reflects the information it represents at a particular position.

For an expanded discussion of the type of raster data supported by ESRI products, review Chapter 9, 'Cell-based modeling with rasters', in *Modeling Our World*.

ArcSDE stores raster datasets similar to the way it stores compressed binary feature classes (see Appendix A, 'ArcSDE compressed binary'). A raster column is added to a business table, and each cell of the raster column contains a reference to a raster stored in a separate raster table. Therefore, each row of a business table references an entire raster.

ArcSDE stores the raster bands in the raster band's table. ArcSDE joins the raster band's table to the raster table on the raster_id column. The raster band table's raster_id column is a foreign key reference to the raster table's raster_id primary key.

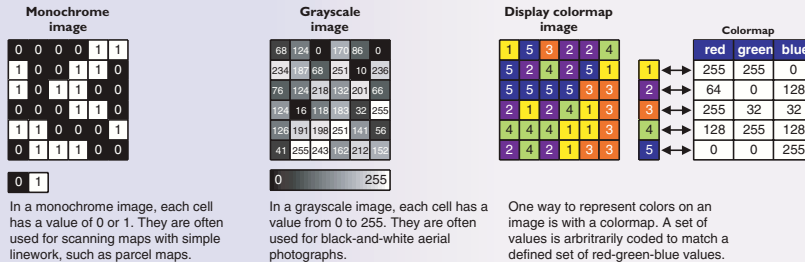
ArcSDE automatically stores any existing image metadata such as image statistics, colormaps, or bit-masks in the raster auxiliary table. The rasterband_id column of the raster auxiliary table is a foreign key reference to the primary key of the raster band's table. ArcSDE joins the two tables on this primary/foreign key reference when accessing a raster band's metadata.

The rendition of rasters

A raster can have one or many bands. The cell values of rasters can be drawn in a variety of ways. These are some of the ways to display rasters by cell values.

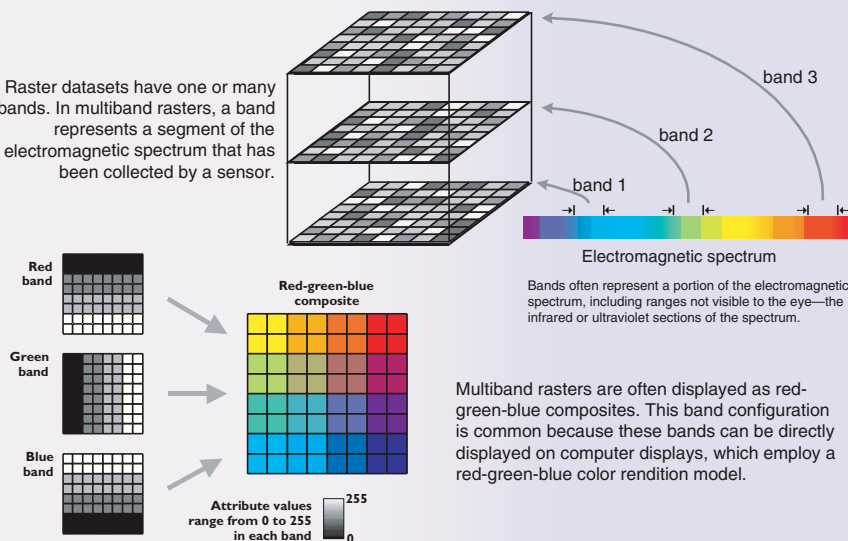
Displaying single-band rasters

Cell values in single-band rasters can be drawn in these three basic ways.



Displaying multiband rasters

Raster datasets have one or many bands. In multiband rasters, a band represents a segment of the electromagnetic spectrum that has been collected by a sensor.



The raster block's table stores the pixels of each raster band. ArcSDE tiles the pixels into blocks according to a user-defined dimension. ArcSDE does not have a default dimension, however, applications that store raster data in ArcSDE do. ArcToolbox and ArcCatalog, for example, use default raster block dimensions of 128-by-128 pixels per block. The dimensions of the raster block along with the compression method, if one is specified, determine the storage size of each raster block. You should select raster block dimensions that, combined

with the compression method, allow each row of the raster block table to fit within SQL Server.


The raster block's table contains the `rasterband_id` column, which is a foreign key reference to the raster band table's `rasterband_id` primary key. ArcSDE joins these tables together on the primary/foreign key reference when accessing the blocks of the raster bands.

ArcSDE populates the raster blocks table according to a declining resolution pyramid. The height of the pyramid is determined by the number of levels, specified by application. The application such as ArcToolbox or ArcCatalog may allow you to define the levels, or it may request that ArcSDE calculate them, or it may offer both possible choices.

The pyramid begins at the base, or level 0, which contains the original pixels of the image. The pyramid proceeds toward the apex by coalescing four pixels from the previous level into a single pixel at the current level. This process continues until less than four pixels remain or until ArcSDE exhausts the defined number of levels.

The apex of the pyramid is reached when the uppermost level has less than four pixels. The addition of the pyramid levels increases the number of raster blocks by one third. However, since it is possible for the user to specify the number of levels, the apex of the pyramid may not be obtainable, and therefore the size increase of the raster blocks table is restricted to the number of resolution levels added.

Figure C.1: When you build a pyramid more rasters are created by progressively downsampling the previous level by a factor of two until the apex. As the application zooms out and the raster cells grow smaller than the resolution threshold, ArcSDE selects a higher level of the pyramid. The purpose of the pyramid is to optimize display performance.



The pyramid allows ArcSDE to provide the application with a constant resolution of pixel data regardless of the rendering window's scale. Data of a large raster transfers quicker to the client when a pyramid exists since ArcSDE can transfer fewer cells of a reduced resolution.

Raster schema

When you import a raster into an ArcSDE database, ArcSDE adds a raster column to the business table of your choice. You may name the raster column whatever you like, so long as

it conforms to SQL Server's column naming convention. ArcSDE restricts one raster column per business table.

The raster column is a foreign key reference to the raster_id column of the raster table created during the addition of the raster column. Also joined to the raster table's raster_id primary key, the raster band table stores the bands of the image. The raster auxiliary table, joined one-to-one to the raster bands table by rasterband_id, stores the metadata of each raster band. The rasterband-id also joins the raster band's table to the raster blocks table in a many-to-one relationship. The raster blocks table rows store blocks of pixels, determined by the dimensions of the block.

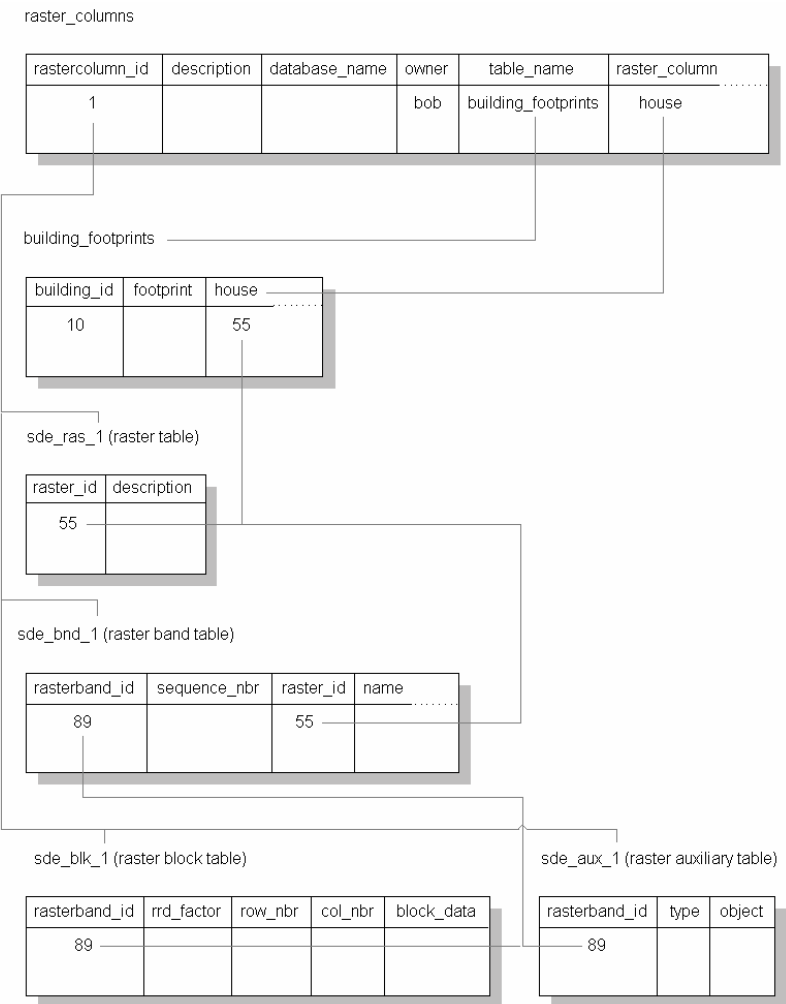


Figure C.2: When ArcSDE adds a raster column to a table, it records that column in the sde user's raster_columns table. The rastercolumn_id is used in the creation of the table names of the raster, raster bands, raster auxiliary, and raster blocks table.

The sections that follow describe the schema of the tables associated with the storage of raster data. Refer to Figure C.2 for an illustration of these tables and the manner in which they are associated with one another.

RASTER_COLUMNS table

When you add a raster column to a business table, ArcSDE adds a record to the RASTER_COLUMNS system table maintained in the sde user's schema. ArcSDE also creates four tables to store the raster images and metadata associated with each one.

NAME	DATA TYPE	NULL?
rastercolumn_id	INT (4)	NOT NULL
description	VARCHAR (65)	NULL
database_name	VARCHAR (32)	NULL
owner	VARCHAR (32)	NOT NULL
table_name	VARCHAR (160)	NOT NULL
raster_column	VARCHAR (32)	NOT NULL
cdate	INT (4)	NOT NULL
config_keyword	VARCHAR (32)	NULL
minimum_id	INT (4)	NULL
base_rastercolumn_id	INT (4)	NOT NULL
rastercolumn_mask	INT (4)	NOT NULL
srid	INT (4)	NULL

Raster columns table

- ⌘ rastercolumn_id (SE_INTEGER_TYPE)—the table's primary key.
- ⌘ description (SE_STRING_TYPE)—The description of the raster table.
- ⌘ database_name (SE_STRING_TYPE)—The name of the database.
- ⌘ owner (SE_STRING_TYPE)—The owner of raster column's business table.
- ⌘ table_name (SE_STRING_TYPE)—The business table name.
- ⌘ raster_column (SE_STRING_TYPE)—The raster column name.
- ⌘ cdate (SE_INTEGER_TYPE)—The date the raster column was added to the business table.
- ⌘ config_keyword (SE_STRING_TYPE)—The DBTUNE configuration keyword whose storage parameters determine how the tables and indexes of the raster are stored in the SQL Server database. For more information on DBTUNE configuration keywords and their storage parameters, review Chapter 3, 'Configuring DBTUNE storage parameters.'
- ⌘ minimum_id (SE_INTEGER_TYPE)—Defined during the creation of the raster, establishes value of the raster table's raster_id column.
- ⌘ base_rastercolumn_id (SE_INTEGER_TYPE)—If a view of the business table is created that includes the raster column, an entry is added to the RASTER_COLUMNS table. The raster column entry of the view will have its own

rastercolumn_id. The base_rastercolumn_id will be the rastercolumn_id of the business table used to create the view. This base_rastercolumn_id maintains referential integrity to the business table. It ensures that actions performed on the business table raster column are reflected in the view. For example, if the business table's raster column is dropped, it will also be dropped from the view (essentially removing the view's raster column entry from the RASTER_COLUMNS table).

- ⌘ rastercolumn_mask (SE_INTEGER_TYPE)—Currently not used, maintained for future use.
- ⌘ rid (SE_INTEGER_TYPE)—The spatial reference ID (srid) is a foreign key reference to the SPATIAL_REFERENCES table. For images that can be georeferenced, the SRID establishes the X and Y offset translation factor and the scale factor for storage of the image coordinates into the 32-bit integer ArcSDE coordinate storage system. It also stores the coordinate reference system the image was created under.

Business table

In the example that follows, the fictitious BUILD_FOOTPRINTS business table contains the raster column house_image. This is a foreign key reference to the raster table created in the user's schema. In this case the raster table contains a record for each raster of a house. It should be noted that images of houses cannot be georeferenced. Therefore, the SRID column of the RASTER_COLUMN record for this raster is NULL.

NAME	DATA TYPE	NULL?
Building_id	INT (4)	NOT NULL
building_footprint	INT (4)	NOT NULL
house_picture	INT (4)	NOT NULL

BUILDING_FOOTPRINTS business table with house image raster column

- ⌘ building_id (SE_INTEGER_TYPE)—the table's primary key
- ⌘ building_footprints (SE_INTEGER_TYPE)—a spatial column and foreign key reference to a feature table containing the building footprints
- ⌘ house_image (SE_INTEGER_TYPE)—a raster column and foreign key reference to a raster table containing the images of the houses located on each building footprint

Raster table (SDE_RAS_<rastercolumn_id>)

The raster table, created as SDE_RAS_<raster_column_id> in the SQL Server, stores a record for each image stored in a raster column. Raster_column_id is assigned by ArcSDE whenever a raster column is created in the database. A record for each raster column in the database is stored in the ArcSDE RASTER_COLUMNS system table maintained in the sde user's schema.

NAME	DATA TYPE	NULL?
Raster_id	INT (4)	NOT NULL
raster_flags	INT (4)	NULL
description	VARCHAR (65)	NULL

Raster description table schema (SDE_RAS_<raster_column_id>)

- ∉ raster_id (SE_INTEGER_TYPE)—The primary key of the raster table and unique sequential identifier of each image stored in the raster table
- ∉ raster_flags (SE_INTEGER_TYPE)—A bitmap set according to the characteristics of stored image
- ∉ description (SE_STRING_TYPE)—A text description of the image (not implemented at ArcSDE 8.1)

Raster band table (SDE_BND_<rastercolumn_id>)

Each image referenced in a raster may be subdivided into one or more raster bands. The raster band table, created as SDE_BND_<rastercolumn_id>, stores the raster bands of each image stored in the raster table. The raster_id column of the raster band table is a foreign key reference to the raster table's raster_id primary key. The rasterband_id column is the raster band table's primary key. Each raster band in the table is uniquely identified by the sequential rasterband_id.

NAME	DATA TYPE	NULL?
rasterband_id	INT (4)	NOT NULL
sequence_nbr	INT (4)	NOT NULL
raster_id	INT (4)	NOT NULL
name	VARCHAR (65)	NULL
band_flags	INT (4)	NOT NULL
band_width	INT (4)	NOT NULL
band_height	INT (4)	NOT NULL
band_types	INT (4)	NOT NULL
block_width	INT (4)	NOT NULL
block_height	INT (4)	NOT NULL
block_origin_x	FLOAT (8)	NOT NULL
block_origin_y	FLOAT (8)	NOT NULL
eminx	FLOAT (8)	NOT NULL
eminy	FLOAT (8)	NOT NULL
emaxx	FLOAT (8)	NOT NULL
emaxy	FLOAT (8)	NOT NULL
cdate	INT (4)	NOT NULL
mdate	INT (4)	NOT NULL

Raster band table schema

- ∄ rasterband_id (SE_INTEGER_TYPE)—The primary key of the raster band table that uniquely identifies each raster band.
- ∄ sequence_nbr (SE_INTEGER_TYPE)—An optional sequential number that can be combined with the raster_id as a composite key as a second way to uniquely identify the raster band.
- ∄ raster_id (SE_INTEGER_TYPE)—The foreign key reference to the raster table's primary key. Uniquely identifies the raster band when combined with sequence_nbr as a composite key.
- ∄ name (SE_STRING_TYPE)—The name of the raster band.
- ∄ band_flags (SE_INTEGER_TYPE)—A bitmap set according to the characteristics of the raster band.
- ∄ band_width (SE_INTEGER_TYPE)—The pixel width of the band.
- ∄ band_height (SE_INTEGER_TYPE)—The pixel height of the band.
- ∄ band_types (SE_INTEGER_TYPE)—A bitmap band compression data.
- ∄ block_width (SE_INTEGER_TYPE)—The pixel width of the band's tiles.

- ⊄ block_height (SE_INTEGER_TYPE)—The pixel height of the band's tiles.
- ⊄ block_origin_x (SE_FLOAT_TYPE)—The leftmost pixel.
- ⊄ block_origin_y (SE_FLOAT_TYPE)—The bottom-most pixel.

If the image has a map extent, the optional eminx, eminy, emaxx, and emaxy will hold the coordinates of the extent.

- ⊄ eminx (SE_FLOAT_TYPE)—The band's minimum x-coordinate.
- ⊄ eminy (SE_FLOAT_TYPE)—The band's minimum y-coordinate.
- ⊄ emaxx (SE_FLOAT_TYPE)—The band's maximum x-coordinate.
- ⊄ emaxy (SE_FLOAT_TYPE)—The band's maximum y-coordinate.
- ⊄ cdate (SE_FLOAT_TYPE)—The creation date.
- ⊄ mdate (SE_FLOAT_TYPE)—The last modification date.

Raster blocks table (SDE_BLK_<rastercolumn_id>)

Created as SDE_BLK_<rastercolumn_id>, the raster blocks table stores the actual pixel data of the raster images. ArcSDE evenly tiles the bands into blocks of pixels. Tiling the raster band data enables efficient storage and retrieval of the raster data. The raster blocks can be configured so that the records of the raster block table fit with a SQL Server data block, avoiding the adverse effects of data block chaining.

The rasterband_id column of the raster block table is a foreign key reference to the raster band table's primary key. A composite unique key is formed by combining the rasterband_id, rrd_factor, row_nbr, and col_nbr columns.

NAME	DATA TYPE	NULL?
rasterband_id	INT (4)	NOT NULL
rrd_factor	INT (4)	NOT NULL
row_nbr	INT (4)	NOT NULL
col_nbr	INT (4)	NOT NULL
block_data	image	NOT NULL

Raster block table schema

- ⊄ rasterband_id (SE_INTEGER_TYPE)—The foreign key reference to the raster band table's primary key.
- ⊄ rrd_factor (SE_INTEGER_TYPE)—The reduced resolution dataset factor determines the position of the raster band block within the resolution pyramid. The resolution pyramid begins at 0 for the highest resolution and increases until the raster's bands lowest resolution level has been reached.
- ⊄ row_nbr (SE_INTEGER_TYPE)—The block's row number.

- col_nbr (SE_INTEGER_TYPE)—The block’s column number.
- block_data (SE_BLOB_TYPE)—The block’s tile of pixel data.

Raster band auxiliary table (SDE_AUX_<rastercolumn_id>)

The raster band auxiliary table, created as SDE_AUX_<rastercolumn_id>, stores optional raster metadata such as the image color map, image statistics, and bit masks used for image overlay and mosaicking. The rasterband_id column is a foreign key reference to the primary key of the raster band table.

NAME	DATA TYPE	NULL?
rasterband_id	INT (4)	NOT NULL
type	INT (4)	NOT NULL
object	image	NOT NULL

Raster auxiliary table schema

- rasterband_id (SE_INTEGER_TYPE)—The foreign key reference to the raster band table’s primary key
- type (SE_INTEGER_TYPE)—A bitmap set according to the characteristics of the data stored in the object column
- object (SE_BLOB_TYPE)—May contain the image color map, image statistics, etc.