


IBM Tivoli Storage Manager
for HP-UX
Version 7.1.7

Installation Guide



IBM Tivoli Storage Manager
for HP-UX
Version 7.1.7

Installation Guide



Note:

Before you use this information and the product it supports, read the information in “Notices” on page 111.

Seventh edition (September 2016)

This edition applies to version 7, release 1, modification 7 of IBM Tivoli Storage Manager (product numbers 5608-E01, 5608-E02, 5608-E03), and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 1993, 2016.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this publication	v
Who should read this guide	v
Installable components	v
Publications	vi

What's new in Version 7.1.7.	vii
-------------------------------------	------------

Part 1. Installing and upgrading the server 1

Chapter 1. Planning to install the server 3

What you should know first	3
Planning for optimal performance	3
Planning for the server hardware and the operating system	4
Planning for the server database disks	8
Planning for the server recovery log disks	11
Planning for directory-container and cloud-container storage pools	12
Planning for storage pools in DISK or FILE device classes	17
Planning for the correct type of storage technology	20
Applying best practices to the server installation	21
Minimum system requirements for the Tivoli Storage Manager server	22
Compatibility of the Tivoli Storage Manager server with other DB2 products on the system	25
IBM Installation Manager	26
Worksheets for planning details for the server	27
Capacity planning	28
Estimating space requirements for the database	28
Recovery log space requirements	31
Monitoring space utilization for the database and recovery logs	44
Deleting installation rollback files	45
Server naming best practices	46
Installation directories	48

Chapter 2. Installing the server components 49

Obtaining the installation package	49
Installing Tivoli Storage Manager by using the installation wizard	50
Installing Tivoli Storage Manager by using console mode	51
Installing Tivoli Storage Manager in silent mode	52
Installing server language packages	53
Server language locales	53
Configuring a language package	54
Updating a language package	55

Chapter 3. Taking the first steps after you install Tivoli Storage Manager . . . 57

Tuning kernel parameters	57
Using the db2osconf utility	58
Updating kernel parameters	58
Suggested minimum values	59
Creating the user ID and directories for the server instance	59
Configuring the Tivoli Storage Manager server	61
Configuring Tivoli Storage Manager by using the configuration wizard	61
Configuring the server instance manually	62
Configuring server options for server database maintenance	70
Starting the server instance	71
Verifying access rights and user limits	71
Starting the server from the instance user ID	73
Automatically starting servers	73
Stopping the server	75
Registering licenses	75
Specifying a device class in preparation for database backups	75
Running multiple server instances on a single system	76
Monitoring the server	77

Chapter 4. Installing a Tivoli Storage Manager server fix pack 79

Chapter 5. Upgrading to V7.1.7 83

Upgrading from V6.2 or V6.3 to V7.1.7	84
Planning the upgrade	84
Preparing the system	85
Installing V7.1.7 and verifying the upgrade	86
Upgrading from V6.1 to V7.1.7	90

Chapter 6. Reverting from Version 7.1 to the previous V6 server. 91

Steps for reverting to the previous server version	91
Additional recovery steps if you created new storage pools or enabled data deduplication	92

Chapter 7. Reference: DB2 commands for Tivoli Storage Manager server databases 95

Chapter 8. Uninstalling Tivoli Storage Manager 99

Uninstalling Tivoli Storage Manager by using a graphical wizard	99
Uninstalling Tivoli Storage Manager in console mode	100
Uninstalling Tivoli Storage Manager in silent mode	100

Uninstalling and reinstalling Tivoli Storage Manager	101
Uninstalling IBM Installation Manager	102

Part 2. Appendixes 103

Appendix A. Installation log files . . . 105

Appendix B. System resource requirements 107

Estimating required semaphore resources	107
Estimating required process resources	107
Estimating required number of threads per process	108

Viewing and modifying the kernel configuration	108
--	-----

Appendix C. Accessibility features for the Tivoli Storage Manager product family. 109

Notices 111

Glossary 115

Index 117

About this publication

This publication contains installation and configuration instructions for the IBM® Tivoli® Storage Manager server, server languages, license, and device driver.

Who should read this guide

This publication is intended for a system administrator installing and configuring a Tivoli Storage Manager server, the Operations Center, or upgrading from Tivoli Storage Manager V6.

If you are upgrading a V7.1 server to a later level of V7.1, see Chapter 4, “Installing a Tivoli Storage Manager server fix pack,” on page 79.

If you are upgrading a V6 server to a newer version, see Chapter 5, “Upgrading to V7.1.7,” on page 83.

Installable components

The IBM Tivoli Storage Manager server and licenses are required components.

Table 1 describes all the installable components.

Table 1. Tivoli Storage Manager installable components

Tivoli Storage Manager component	Description	Additional information
Server (required)	Includes the database, the Global Security Kit (GSKit), IBM Java™ Runtime Environment (JRE), and tools to help you configure and manage the server.	See Chapter 2, “Installing the server components,” on page 49.
Language package (optional)	Each language package (one for each language) contains language-specific information for the server.	See “Installing server language packages” on page 53.
Licenses (required)	Includes support for all licensed features. After you install this package, you must register the licenses you purchased.	Use the REGISTER LICENSE command.
Devices (optional)	Extends media management capability.	A list of devices that are supported by this driver is available from the IBM Support Portal.

Table 1. Tivoli Storage Manager installable components (continued)

Tivoli Storage Manager component	Description	Additional information
Storage agent (optional)	<p>Installs the component that allows client systems to write data directly to, or read data directly from, storage devices that are attached to a storage area network (SAN).</p> <p>Remember: IBM Tivoli Storage Manager for Storage Area Networks is a separately licensed product.</p>	<p>For more information about storage agents, see Tivoli Storage Manager for Storage Area Networks (V7.1.1).</p>

Publications

The Tivoli Storage Manager product family includes IBM Tivoli Storage FlashCopy[®] Manager, IBM Tivoli Storage Manager for Space Management, IBM Tivoli Storage Manager for Databases, and several other storage management products from IBM.

To view IBM product documentation, see IBM Knowledge Center.

What's new in Version 7.1.7

IBM Tivoli Storage Manager Version 7.1.7 introduces new features and updates.

For a list of new features and updates in this release and previous Version 7 releases, see What's new.

New and changed information in this product documentation is indicated by a vertical bar (!) to the left of the change.

Part 1. Installing and upgrading the server

Install and upgrade the Tivoli Storage Manager server.

Chapter 1. Planning to install the server

Install the server software on the computer that manages storage devices and install the client software on every workstation that transfers data to Tivoli Storage Manager server-managed storage.

What you should know first

Before installing IBM Tivoli Storage Manager, be familiar with your operating systems, storage devices, communication protocols, and system configurations.

Server maintenance releases, client software, and publications are available from the IBM Support Portal.

Restriction: You can install and run the Version 7.1.7 server on a system that already has DB2® installed on it, whether DB2 was installed independently or as part of some other application, with some restrictions. For details, see “Compatibility of the Tivoli Storage Manager server with other DB2 products on the system” on page 25.

Experienced DB2 administrators can choose to perform advanced SQL queries and use DB2 tools to monitor the database. Do not, however, use DB2 tools to change DB2 configuration settings from those that are preset by Tivoli Storage Manager, or alter the DB2 environment for Tivoli Storage Manager in other ways, such as with other products. The V7.1.7 server has been built and tested extensively using the data definition language (DDL) and database configuration that the server deploys.

Attention: Do not alter the DB2 software that is installed with Tivoli Storage Manager installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack of DB2 software because doing so can damage the database.

Planning for optimal performance

Before you install the Tivoli Storage Manager server, evaluate the characteristics and configuration of the system to ensure that the server is set up for optimal performance.

Procedure

1. Review “What you should know first.”
2. Review each of the following sub-sections.

Planning for the server hardware and the operating system

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
<p>Does the operating system and hardware meet or exceed requirements?</p> <ul style="list-style-type: none"> • Number and speed of processors • System memory • Supported operating system level 	<p>If you are using the minimum required amount of memory, you can support a minimal workload.</p> <p>You can experiment by adding more system memory to determine whether the performance is improved. Then, decide whether you want to keep the system memory dedicated to the server. Test the memory variations by using the entire daily cycle of the server workload.</p> <p>If you run multiple servers on the system, add the requirements for each server to get the requirements for the system.</p> <p>Restriction: For servers on AIX® operating systems, do not use Active Memory™ Expansion (AME). When you use AME, the IBM DB2 software uses 4 KB pages instead of 64 KB pages. Each 4 KB page must be decompressed when accessed, and compressed when not needed. When the compression or decompression occurs, DB2 and the server wait for access to the page, which degrades the server performance.</p>	<p>Review operating system requirements at Tivoli Storage Manager Supported Operating Systems.</p> <p>Additionally, review the guidance in Tuning tasks for operating systems and other applications.</p> <p>For more information about requirements when these features are in use, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication <p>For more information about sizing requirements for the server and storage, see the Tivoli Storage Manager Blueprint.</p>
<p>Are disks configured for optimal performance?</p>	<p>The amount of tuning that can be done for different disk systems varies. Ensure that the appropriate queue depths and other disk system options are set.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • "Planning for server database disks" • "Planning for server recovery log disks" • "Planning for storage pools in DISK or FILE device classes"

Question	Tasks, characteristics, options, or settings	More information
Does the server have enough memory?	<p>Heavier workloads and advanced features such as data deduplication and node replication require more than the minimum system memory that is specified in the system requirements document.</p> <p>For databases that are not enabled for data deduplication, use the following guidelines to specify memory requirements:</p> <ul style="list-style-type: none"> • For databases less than 500 GB, you need 16 GB of memory. • For databases with a size of 500 GB - 1 TB, you need 24 GB of memory. • For databases with a size of 1 TB - 1.5 TB, you need 32 GB of memory. • For databases greater than 1.5 TB, you need 40 GB of memory. <p>Ensure that you allocate extra space for the active log and the archive log for replication processing.</p>	<p>For more information about requirements when these features are in use, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication • Memory requirements
Does the system have enough host bus adapters (HBAs) to handle the data operations that the Tivoli Storage Manager server must run simultaneously?	<p>Understand what operations require use of HBAs at the same time.</p> <p>For example, a server must store 1 GB/sec of backup data while also doing storage pool migration that requires 0.5 GB/sec capacity to complete. The HBAs must be able to handle all of the data at the speed required.</p>	See Tuning HBA capacity.

Installing the Tivoli Storage Manager server

Question	Tasks, characteristics, options, or settings	More information
Is network bandwidth greater than the planned maximum throughput for backups?	<p>Network bandwidth must allow the system to complete operations such as backups in the time that is allowed or that meets service level commitments.</p> <p>For node replication, network bandwidth must be greater than the planned maximum throughput.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none">• Tuning network performance• Checklist for node replication

Question	Tasks, characteristics, options, or settings	More information
<p>Are you using a preferred file system for Tivoli Storage Manager server files?</p>	<p>Use a file system that ensures optimal performance and data availability. The server uses direct I/O with file systems that support the feature. Using direct I/O can improve throughput and reduce processor use. The following list identifies the preferred file systems for each operating system:</p> <ul style="list-style-type: none"> • AIX: Use the JFS2 file system with the rbrw option. • HP-UX: Use the VxFS file system. • Linux: Use either the ext3 or ext4 file system for the database, recovery log, and storage pool data. Use the following file system that is appropriate for your operating system and level: <ul style="list-style-type: none"> – For Red Hat Enterprise Linux x86_64, use the ext3 or ext4 file system. If Red Hat Enterprise Linux 6.4 or later is installed, use the ext4 file system. – For SUSE Linux Enterprise Server and for Red Hat Enterprise Linux ppc64, use the ext3 file system. • Oracle Solaris: Use the ZFS file system. • Windows: Use New Technology File System (NTFS) without compression. 	<p>For more information, see Configuring the operating system for disk performance.</p>

Installing the Tivoli Storage Manager server

Question	Tasks, characteristics, options, or settings	More information
Are you planning to configure enough paging space?	<p>Paging space, or swap space, extends the memory that is available for processing. When the amount of free RAM in the system is low, programs or data that is not in use are moved from memory to paging space. This action releases memory for other activities, such as database operations.</p> <p>For AIX, Linux, HP-UX, and Oracle Solaris systems, use a minimum of 32 GB of paging space or 50% of your RAM, whichever value is larger.</p> <p>On Windows systems, paging space is automatically configured.</p>	
Are you planning to tune the kernel parameters after installation of the server?	You must tune kernel parameters for servers that are on HP-UX, Linux, or Oracle Solaris systems.	<p>See the information about tuning kernel parameters for your operating system:</p> <ul style="list-style-type: none">• HP-UX: Tuning kernel parameters for HP-UX systems• Linux: Tuning kernel parameters for Linux systems• Solaris: Tuning kernel parameters for Oracle Solaris systems <p>Not all operating systems require that kernel parameters be tuned.</p>

Planning for the server database disks

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
Is the database on fast, low-latency disks?	<p>Do not use the following drives for the Tivoli Storage Manager database:</p> <ul style="list-style-type: none"> Nearline SAS (NL-SAS) Serial Advanced Technology Attachment (SATA) Parallel Advanced Technology Attachment (PATA) <p>Do not use internal disks that are included by default in most server hardware.</p> <p>Enterprise-grade solid-state disks (SSD), with Fibre Channel or SAS interface, offer the best performance.</p> <p>If you plan to use the data deduplication functions of Tivoli Storage Manager, focus on disk performance in terms of I/O operations per second (IOPS).</p>	For more information, see Checklist for data deduplication.
Is the database stored on disks or LUNs that are separate from disks or LUNs that are used for the active log, archive log, and storage pool volumes?	<p>Separation of the server database from other server components helps reduce contention for the same resources by different operations that must run at the same time.</p> <p>Tip: The database and the archive log can share an array when you use solid-state drive (SSD) technology.</p>	
If you are using RAID, do you know how to select the optimal RAID level for your system? Are you defining all LUNs with the same size and type of RAID?	<p>When a system must do large numbers of writes, RAID 10 outperforms RAID 5. However, RAID 10 requires more disks than RAID 5 for the same amount of usable storage.</p> <p>If your disk system is RAID, define all your LUNs with the same size and type of RAID. For example, do not mix 4+1 RAID 5 with 4+2 RAID 6.</p>	
If an option to set the strip size or segment size is available, are you planning to optimize the size when you configure the disk system?	If you can set the strip size or segment size, use 64 KB or 128 KB sizes on disk systems for the database.	The block size that is used for the database varies depending on the table space. Most table spaces use 8 KB blocks, but some use 32 KB blocks.

Installing the Tivoli Storage Manager server

Question	Tasks, characteristics, options, or settings	More information
<p>Are you planning to create at least four directories, also called storage paths, on four separate LUNs for the database?</p> <p>Create one directory per distinct array on the subsystem. If you have fewer than three arrays, create a separate LUN volume within the array.</p>	<p>Heavier workloads and use of some features require more database storage paths than the minimum requirements.</p> <p>Server operations such as data deduplication drive a high number of input/output operations per second (IOPS) for the database. Such operations perform better when the database has more directories.</p> <p>For server databases that are larger than 2 TB or are expected to grow to that size, use eight directories.</p> <p>Consider planned growth of the system when you determine how many storage paths to create. The server uses the higher number of storage paths more effectively if the storage paths are present when the server is first created.</p> <p>Use the <code>DB2_PARALLEL_IO</code> variable to force parallel I/O to occur on table spaces that have one container, or on table spaces that have containers on more than one physical disk. If you do not set the <code>DB2_PARALLEL_IO</code> variable, I/O parallelism is equal to the number of containers that are used by the table space. For example, if a table space spans four containers, the level of I/O parallelism that is used is 4.</p>	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • Checklist for data deduplication • Checklist for node replication <p>For help with forecasting growth when the server deduplicates data, see technote 1596944.</p> <p>For the most recent information about database size, database reorganization, and performance considerations for Tivoli Storage Manager Version 6 and Version 7 servers, see technote 1452146.</p> <p>For information about setting the <code>DB2_PARALLEL_IO</code> variable, see Recommended settings for IBM DB2 registry variables.</p>
<p>Are all directories for the database the same size?</p>	<p>Directories that are all the same size ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching.</p> <p>This guideline also applies if you must add storage paths after the initial configuration of the server.</p>	
<p>Are you planning to raise the queue depth of the database LUNs on AIX systems?</p>	<p>The default queue depth is often too low.</p>	<p>See Configuring AIX systems for disk performance.</p>

Planning for the server recovery log disks

Use the checklist to verify that the system where the server is installed meets requirements for hardware and software configuration.

Question	Tasks, characteristics, options, or settings	More information
Are the active log and archive log stored on disks or LUNs that are separate from what is used for the database and storage pool volumes?	Ensure that the disks where you place the active log are not used for other server or system purposes. Do not place the active log on disks that contain the server database, the archive log, or system files such as page or swap space.	Separation of the server database, active log, and archive log helps to reduce contention for the same resources by different operations that must run at the same time.
Are the logs on disks that have nonvolatile write cache?	Nonvolatile write cache allows data to be written to the logs as fast as possible. Faster write operations for the logs can improve performance for server operations.	
Are you setting the logs to a size that adequately supports the workload?	<p>If you are not sure about the workload, use the largest size that you can.</p> <p>Active log The maximum size is 512 GB, set with the ACTIVELOGSIZE server option.</p> <p>Ensure that there is at least 8 GB of free space on the active log file system after the fixed size active logs are created.</p> <p>Archive log The size of the archive log is limited by the size of the file system on which it is located, and not by a server option. Make the archive log at least as large as the active log.</p>	<ul style="list-style-type: none"> For log sizing details, see the recovery log information in technote 1421060. For information about sizing when you use data deduplication, see Checklist for data deduplication.
Are you defining an archive failover log? Are you placing this log on a disk that is separate from the archive log?	The archive failover log is for emergency use by the server when the archive log becomes full. Slower disks can be used for the archive failover log.	<p>Use the ARCHFAILOVERLOGDIRECTORY server option to specify the location of the archive failover log.</p> <p>Monitor the usage of the directory for the archive failover log. If the archive failover log must be used by the server, the space for the archive log might not be large enough.</p>

Installing the Tivoli Storage Manager server

Question	Tasks, characteristics, options, or settings	More information
If you are mirroring the active log, are you using only one type of mirroring?	<p>You can mirror the log by using one of the following methods. Use only one type of mirroring for the log.</p> <ul style="list-style-type: none">• Use the MIRRORLOGDIRECTORY option that is available for the Tivoli Storage Manager server to specify a mirror location.• Use software mirroring, such as Logical Volume Manager (LVM) on AIX.• Use mirroring in the disk system hardware.	<p>If you mirror the active log, ensure that the disks for both the active log and the mirror copy have equal speed and reliability.</p> <p>For more information, see <i>Configuring and tuning the recovery log</i>.</p>

Planning for directory-container and cloud-container storage pools

Review how your directory-container and cloud-container storage pools are set up to ensure optimal performance.

Question	Tasks, characteristics, options, or settings	More information
Measured in terms of input/output operations per second (IOPS), are you using fast disk storage for the Tivoli Storage Manager database?	<p>Use a high-performance disk for the database. Use solid-state drive technology for data deduplication processing.</p> <p>Ensure that the database has a minimum capability of 3000 IOPS. For each TB of data that is backed up daily (before data deduplication), add 1000 IOPS to this minimum.</p> <p>For example, a Tivoli Storage Manager server that is ingesting 3 TB of data per day would need 6000 IOPS for the database disks:</p> $3000 \text{ IOPS minimum} + 3000 (3 \text{ TB} \times 1000 \text{ IOPS}) = 6000 \text{ IOPS}$	<p>For recommendations about disk selection, see "Planning for server database disks".</p> <p>For more information about IOPS, see the Tivoli Storage Manager Blueprints.</p>

Question	Tasks, characteristics, options, or settings	More information
Do you have enough memory for the size of your database?	<p>Use a minimum of 40 GB of system memory for Tivoli Storage Manager servers, with a database size of 100 GB, that are deduplicating data. If the retained capacity of backup data grows, the memory requirement might need to be higher.</p> <p>Monitor memory usage regularly to determine whether more memory is required.</p> <p>Use more system memory to improve caching of database pages. The following memory size guidelines are based on the daily amount of new data that you back up:</p> <ul style="list-style-type: none"> • 128 GB of system memory for daily backups of data, where the database size is 1 - 2 TB • 192 GB of system memory for daily backups of data, where the database size is 2 - 4 TB 	Memory requirements
Have you properly sized the storage capacity for the database active log and archive log?	<p>Configure the server to have a minimum active log size of 128 GB by setting the ACTIVELOGSIZE server option to a value of 131072.</p> <p>The suggested starting size for the archive log is 1 TB. The size of the archive log is limited by the size of the file system on which it is located, and not by a server option. Ensure that there is at least 10% extra disk space for the file system than the size of the archive log.</p> <p>Use a directory for the database archive logs with an initial free capacity of at least 1 TB. Specify the directory by using the ARCHLOGDIRECTORY server option.</p> <p>Define space for the archive failover log by using the ARCHFAILOVERLOGDIRECTORY server option.</p>	For more information about sizing for your system, see the Tivoli Storage Manager Blueprints.

Installing the Tivoli Storage Manager server

Question	Tasks, characteristics, options, or settings	More information
Is compression enabled for the archive log and database backups?	<p>Enable the ARCHLOGCOMPRESS server option to save storage space.</p> <p>This compression option is different from inline compression. Inline compression is enabled by default with Tivoli Storage Manager V7.1.5 and later.</p> <p>Restriction: Do not use this option if the amount of backed up data exceeds 6 TB per day.</p>	For more information about compression for your system, see the Tivoli Storage Manager Blueprints.
<p>Are the Tivoli Storage Manager database and logs on separate disk volumes (LUNs)?</p> <p>Is the disk that is used for the database configured according to best practices for a transactional database?</p>	The database must not share disk volumes with Tivoli Storage Manager database logs or storage pools, or with any other application or file system.	For more information about server database and recovery log configuration, see Server database and recovery log configuration and tuning.
Are you using a minimum of eight (2.2 GHz or equivalent) processor cores for each Tivoli Storage Manager server that you plan to use with data deduplication?	If you are planning to use client-side data deduplication, verify that client systems have adequate resources available during a backup operation to complete data deduplication processing. Use a processor that is at least the minimum equivalent of one 2.2 GHz processor core per backup process with client-side data deduplication.	<ul style="list-style-type: none"> • Effective planning and use of deduplication • Tivoli Storage Manager Blueprints
Did you allocate enough storage space for the database?	<p>For a rough estimate, plan for 100 GB of database storage for every 50 TB of data that is to be protected in deduplicated storage pools. <i>Protected data</i> is the amount of data before data deduplication, including all versions of objects stored.</p> <p>As a best practice, define a new container storage pool exclusively for data deduplication. Data deduplication occurs at the storage-pool level, and all data within a storage pool, except encrypted data, is deduplicated.</p>	

Question	Tasks, characteristics, options, or settings	More information
Have you estimated storage pool capacity to configure enough space for the size of your environment?	<p>You can estimate capacity requirements for a deduplicated storage pool by using the following technique:</p> <ol style="list-style-type: none"> 1. Estimate the base size of the source data. 2. Estimate the daily backup size by using an estimated change and growth rate. 3. Determine retention requirements. 4. Estimate the total amount of source data by factoring in the base size, daily backup size, and retention requirements. 5. Apply the deduplication ratio factor. 6. Apply the compression ratio factor. 7. Round up the estimate to consider transient storage pool usage. 	For an example of using this technique, see Effective planning and use of deduplication.
Have you distributed disk I/O over many disk devices and controllers?	<p>Use arrays that consist of as many disks as possible, which is sometimes referred to as wide striping. Ensure that you use one database directory per distinct array on the subsystem.</p> <p>Set the <i>DB2_PARALLEL_IO</i> registry variable to enable parallel I/O for each table space used if the containers in the table space span multiple physical disks.</p> <p>When I/O bandwidth is available and the files are large, for example 1 MB, the process of finding duplicates can occupy the resources of an entire processor. When files are smaller, other bottlenecks can occur.</p> <p>Specify eight or more file systems for the deduplicated storage pool device class so that I/O is distributed across as many LUNs and physical devices as possible.</p>	<p>For guidelines about setting up storage pools, see "Planning for storage pools in DISK or FILE device classes".</p> <p>For information about setting the <i>DB2_PARALLEL_IO</i> variable, see Recommended settings for IBM DB2 registry variables.</p>
Have you scheduled daily operations based on your backup strategy?	<p>The best practice sequence of operations is in the following order:</p> <ol style="list-style-type: none"> 1. Client backup 2. Storage pool protection 3. Node replication 4. Database backup 5. Expire inventory 	<ul style="list-style-type: none"> • Scheduling data deduplication and node replication processes • Daily operations for directory-container storage pools

Installing the Tivoli Storage Manager server

Question	Tasks, characteristics, options, or settings	More information
Do you have enough storage to manage the DB2 lock list?	<p>If you deduplicate data that includes large files or large numbers of files concurrently, the process can result in insufficient storage space. When the lock list storage is insufficient, backup failures, data management process failures, or server outages can occur.</p> <p>File sizes greater than 500 GB that are processed by data deduplication are most likely to deplete storage space. However, if many backup operations use client-side data deduplication, this problem can also occur with smaller-sized files.</p>	For information about tuning the DB2 LOCKLIST parameter, see Tuning server-side data deduplication.
Is sufficient bandwidth available to transfer data to a Tivoli Storage Manager server?	<p>To transfer data to a Tivoli Storage Manager server, use client-side or server-side data deduplication and compression to reduce the bandwidth that is required.</p> <p>Use a V7.1.5 server or higher to use inline compression and use a V7.1.6 or later client to enable enhanced compression processing.</p>	For more information, see the enablededup client option.
Have you determined how many storage pool directories to assign to each storage pool?	<p>Assign directories to a storage pool by using the DEFINE STGPOOLDIRECTORY command.</p> <p>Create multiple storage pool directories and ensure that each directory is backed up to a separate disk volume (LUN).</p>	
Did you allocate enough disk space in the cloud-container storage pool?	<p>To prevent backup failures, ensure that the local directory has enough space. Use the following list as a guide for optimal disk space:</p> <ul style="list-style-type: none"> • For serial-attached SCSI (SAS) and spinning disk, calculate the amount of new data that is expected after daily data reduction (compression and data deduplication). Allocate up to 100 percent of that amount, in terabytes, for disk space. • Provide 3 TB for flash-based storage systems with fast network connections to on-premises, high-performance cloud systems. • Provide 5 TB for solid-state drive (SSD) systems with fast network connections to high-performance cloud systems. 	

Question	Tasks, characteristics, options, or settings	More information
Did you select the appropriate type of local storage?	<p>Ensure that data transfers from local storage to cloud finish before the next backup cycle starts.</p> <p>Tip: Data is removed from local storage soon after it moves to the cloud.</p> <p>Use the following guidelines:</p> <ul style="list-style-type: none"> • Use flash or SSD for large systems that have high-performing cloud systems. Ensure that you have a dedicated 10 GB wide area network (WAN) link with a high-speed connection to the object storage. For example, use flash or SSD if you have a dedicated 10 GB WAN link plus a high-speed connection to either a Cleversafe® location or to an Amazon Simple Storage Service (Amazon S3) data center. • Use larger capacity 15000 rpm SAS disks for these scenarios: <ul style="list-style-type: none"> – Medium-sized systems – Slower cloud connections, for example, 1 GB – When you use Cleversafe as your service provider across several regions • For SAS or spinning disk, calculate the amount of new data that is expected after daily data reduction (compression and data deduplication). Allocate up to 100 percent of that amount for disk space, in terabytes. 	

Planning for storage pools in DISK or FILE device classes

Use the checklist to review how your disk storage pools are set up. This checklist includes tips for storage pools that use DISK or FILE device classes.

Installing the Tivoli Storage Manager server

Question	Tasks, characteristics, options, or settings	More information
Can the storage pool LUNs sustain throughput rates for 256 KB sequential reads and writes to adequately handle the workload within the time constraints?	<p>When you are planning for peak loads, consider all the data that you want the server to read or write to the disk storage pools simultaneously. For example, consider the peak flow of data from client backup operations and server data-movement operations such as migration that run at the same time.</p> <p>The Tivoli Storage Manager server reads and writes to storage pools predominantly in 256 KB blocks.</p> <p>If the disk system includes the capability, configure the disk system for optimal performance with sequential read/write operations rather than random read/write operations.</p>	For more information, see Analyzing the basic performance of disk systems.
Is the disk configured to use read and write cache?	Use more cache for better performance.	
For storage pools that use FILE device classes, have you determined a good size to use for the storage pool volumes?	Review the information in Optimal number and size of volumes for storage pools that use disk. If you do not have the information to estimate a size for FILE device class volumes, start with volumes that are 50 GB.	Typically, problems arise more frequently when the volumes are too small. Few problems are reported when volumes are larger than needed. When you determine the volume size to use, as a precaution choose a size that might be larger than necessary.
For storage pools that use FILE device classes, are you using preallocated volumes?	<p>Scratch volumes can cause file fragmentation.</p> <p>To ensure that a storage pool does not run out of volumes, set the MAXSCRATCH parameter to a value greater than zero.</p>	<p>Use the DEFINE VOLUME server command to preallocate volumes in the storage pool.</p> <p>Use the DEFINE STGPOOL or UPDATE STGPOOL server command to set the MAXSCRATCH parameter.</p>
For storage pools that use FILE device classes, have you compared the maximum number of client sessions to the number of volumes that are defined?	Always maintain enough usable volumes in the storage pools to allow for the expected peak number of client sessions that run at one time. The volumes might be scratch volumes, empty volumes, or partly filled volumes.	For storage pools that use FILE device classes, only one session or process can write to a volume at the same time.

Question	Tasks, characteristics, options, or settings	More information
<p>For storage pools that use FILE device classes, have you set the MOUNTLIMIT parameter of the device class to a value that is high enough to account for the number of volumes that might be mounted in parallel?</p>	<p>For storage pools that use data deduplication, the MOUNTLIMIT parameter is typically in the range of 500 - 1000.</p> <p>Set the value for MOUNTLIMIT to the maximum number of mount points that are needed for all active sessions. Consider parameters that affect the maximum number of mount points that are needed:</p> <ul style="list-style-type: none"> • The MAXSESSIONS server option, which is the maximum number of Tivoli Storage Manager sessions that can run concurrently. • The MAXNUMMP parameter, which sets the maximum number of mount points that each client node can use. <p>For example, if the maximum number of client node backup sessions is typically 100 and each of the nodes has MAXNUMMP=2, multiply 100 nodes by the 2 mount points for each node to get the value of 200 for the MOUNTLIMIT parameter.</p>	<p>Use the REGISTER NODE or UPDATE NODE server command to set the MAXNUMMP parameter for client nodes.</p>
<p>For storage pools that use DISK device classes, have you determined how many storage pool volumes to put on each file system?</p>	<p>How you configure the storage for a storage pool that uses a DISK device class depends on whether you are using RAID for the disk system.</p> <p>If you are not using RAID, then configure one file system per physical disk, and define one storage pool volume for each file system.</p> <p>If you are using RAID 5 with $n + 1$ volumes, configure the storage in one of the following ways:</p> <ul style="list-style-type: none"> • Configure n file systems on the LUN and define one storage pool volume per file system. • Configure one file system and n storage pool volumes for the LUN. 	<p>For an example layout that follows this guideline, see Sample layout of server storage pools.</p>
<p>Did you create your storage pools to distribute I/O across multiple file systems?</p>	<p>Ensure that each file system is on a different LUN on the disk system.</p> <p>Typically, having 10 - 30 file systems is a good goal, but ensure that the file systems are no smaller than approximately 250 GB.</p>	<p>For details, see the following topics:</p> <ul style="list-style-type: none"> • Tuning disk storage for the server • Tuning and configuring storage pools and volumes

Planning for the correct type of storage technology

Storage devices have different capacity and performance characteristics. These characteristics affect which devices are better for use with Tivoli Storage Manager.

Procedure

Review the following table to help you to choose the correct type of storage technology for the storage resources that the server requires.

Table 2. Storage technology types for Tivoli Storage Manager storage requirements

Storage technology type	Database	Active log	Archive log and archive failover log	Storage pools
Solid-state disk (SSD)	Place the database on SSD in the following circumstances: <ul style="list-style-type: none"> You are using Tivoli Storage Manager data deduplication. You are backing up more than 8 TB of new data daily. 	If you place the Tivoli Storage Manager database on an SSD, as a best practice, place the active log on an SSD. If space is not available, use high-performance disk instead.	Save SSDs for use with the database and active log. The archive log and archive failover logs can be placed on slower storage technology types.	Save SSDs for use with the database and active log. Storage pools can be placed on slower storage technology types.
High-performance disk with the following characteristics: <ul style="list-style-type: none"> 15k rpm disk Fibre Channel or serial-attached SCSI (SAS) interface 	Use high-performance disks in the following circumstances: <ul style="list-style-type: none"> The server does not do data deduplication. The server does not do node replication. Isolate the server database from its logs and storage pools, and from data for other applications.	Use high-performance disks in the following circumstances: <ul style="list-style-type: none"> The server does not do data deduplication. The server does not do node replication. For performance and availability, isolate the active log from the server database, archive logs, and storage pools.	You can use high-performance disks for the archive log and archive failover logs. For availability, isolate these logs from the database and active log.	Use high-performance disks for storage pools in the following circumstances: <ul style="list-style-type: none"> Data is frequently read. Data is frequently written. For performance and availability, isolate storage pool data from the server database and logs, and from data for other applications.
Medium-performance or high-performance disk with the following characteristics: <ul style="list-style-type: none"> 10k rpm disk Fibre Channel or SAS interface 	If the disk system has a mix of disk technologies, use the faster disks for the database and active log. Isolate the server database from its logs and storage pools, and from data for other applications.	If the disk system has a mix of disk technologies, use the faster disks for the database and active log. For performance and availability, isolate the active log from the server database, archive logs, and storage pools.	You can use medium-performance or high-performance disk for the archive log and archive failover logs. For availability, isolate these logs from the database and active log.	Use medium-performance or high-performance disk for storage pools in the following circumstances: <ul style="list-style-type: none"> Data is frequently read. Data is frequently written. For performance and availability, isolate storage pool data from the server database and logs, and from data for other applications.

Table 2. Storage technology types for Tivoli Storage Manager storage requirements (continued)

Storage technology type	Database	Active log	Archive log and archive failover log	Storage pools
SATA, network-attached storage	Do not use this storage for the database. Do not place the database on XIV storage systems.	Do not use this storage for the active log.	Use of this slower storage technology is acceptable because these logs are written once and infrequently read.	Use this slower storage technology in the following circumstances: <ul style="list-style-type: none"> • Data is infrequently written, for example written once. • Data is infrequently read.
Tape and virtual tape				Use for long-term retention or if data is infrequently used.

Applying best practices to the server installation

Typically, hardware configuration and selection have the most significant effect on the performance of a Tivoli Storage Manager solution. Other factors that affect performance are the operating system selection and configuration, and the configuration of Tivoli Storage Manager.

Procedure

- The following best practices are the most important for optimal performance and problem prevention.
- Review the table to determine the best practices that apply to your environment.

Best practice	More information
Use fast disks for the server database. Enterprise-grade solid-state disks (SSD), with Fibre Channel or SAS interface, offer the best performance.	Use fast, low-latency disks for the database. Using SSD is essential if you are using data deduplication and node replication. Avoid Serial Advanced Technology Attachment (SATA) and Parallel Advanced Technology Attachment (PATA) disks. For details and more tips, see the following topics: <ul style="list-style-type: none"> • "Planning for server database disks" • "Planning for the correct type of storage technology"
Ensure that the server system has enough memory.	Review operating system requirements in technote 1243309. Heavier workloads require more than the minimum requirements. Advanced features such as data deduplication and node replication can require more than the minimum memory that is specified in the system requirements document. <p>If you plan to run multiple instances, each instance requires the memory that is listed for one server. Multiply the memory for one server by the number of instances that are planned for the system.</p>

Installing the Tivoli Storage Manager server

Best practice	More information
Separate the server database, the active log, the archive log, and disk storage pools from each other.	Keep all Tivoli Storage Manager storage resources on separate disks. Keep storage pool disks separate from the disks for the server database and logs. Storage pool operations can interfere with database operations when both are on the same disks. Ideally, the server database and logs are also separated from each other. For details and more tips, see the following topics: <ul style="list-style-type: none">• "Planning for server database disks"• "Planning for server recovery log disks"• "Planning for storage pools in DISK or FILE device classes"
Use at least four directories for the server database. For larger servers or servers that use advanced features, use eight directories.	Place each directory on a LUN that is isolated from other LUNs and from other applications. A server is considered to be large if its database is larger than 2 TB or is expected to grow to that size. Use eight directories for such servers. See "Planning for server database disks".
If you are using data deduplication, node replication, or both, follow the guidelines for database configuration and other items.	Configure the server database according to the guidelines, because the database is extremely important to how well the server runs when these features are being used. For details and more tips, see the following topics: <ul style="list-style-type: none">• Checklist for data deduplication• Checklist for node replication
For storage pools that use FILE type device classes, follow the guidelines for the size of storage pool volumes. Typically, 50 GB volumes are best.	Review the information in Optimal number and size of volumes for storage pools that use disk to help you to determine volume size. Configure storage pool devices and file systems based on throughput requirements, not only on capacity requirements. Isolate the storage devices that are used by Tivoli Storage Manager from other applications that have high I/O, and ensure that there is enough throughput to that storage. For more details, see Checklist for storage pools on DISK or FILE.
Schedule Tivoli Storage Manager client operations and server maintenance activities to avoid or minimize overlap of operations.	For more details, see the following topics: <ul style="list-style-type: none">• Tuning the schedule for daily operations• Checklist for server configuration
Monitor operations constantly.	By monitoring, you can find problems early and more easily identify causes. Keep records of monitoring reports for up to a year to help you identify trends and plan for growth. See Monitoring and maintaining the environment for performance.

Minimum system requirements for the Tivoli Storage Manager server

To install Tivoli Storage Manager server on an HP-UX system, it is necessary to have a minimum level of hardware and software, including a communication method and the most current device driver.

Hardware requirements

These tables list the minimum hardware and software requirements for the installation of a Tivoli Storage Manager server. Use these requirements as a starting point. You can find the most current information about system requirements at Tivoli Storage Manager Supported Operating Systems.

Table 3 describes the minimum hardware requirements that are needed for a server on an HP-UX system. The installation fails if you do not have the minimum requirements. For more details about planning disk space, see “Capacity planning” on page 28.

Table 3. Hardware requirements

Type of hardware	Hardware requirements
Hardware	A 64-bit Intel Itanium system.
Disk space	<p>The following minimum values for disk space:</p> <ul style="list-style-type: none"> • The /var directory space requirements for new installations and version upgrades: <ul style="list-style-type: none"> – 512 MB for new installations – 2560 MB for version upgrades • 6 GB in the installation directory • 2 GB for the /opt directory if you do not create mount points • 2 GB for the /tmp directory • 2 GB in the home directory <p>Tip: Expect to use more space for problem determination.</p> <ul style="list-style-type: none"> • 2 GB for the shared resources area <p>Significant additional disk space is required for database and log files. The size of the database depends on the number of client files to be stored and the method by which the server manages them. The default active log space is 16 GB, the minimum that is needed for most workloads and configurations. When you create the active log, you need at least 64 GB in size to run replication. If replication and data deduplication are both being used, create an active log of 128 GB in size. Allocate at least three times the default active log space for the archive log (48 GB). Ensure that you have sufficient resources if you are using data deduplication or expect a heavy client workload.</p> <p>For optimal performance and to facilitate I/O, specify at least two equally sized containers or Logical Unit Numbers (LUNs) for the database. In addition, each active log and archive log should have its own container or LUN.</p> <p>Ensure that you see the capacity planning section for more details about disk space.</p>
Memory	<p>The following minimum values for memory:</p> <ul style="list-style-type: none"> • 12 GB. • 16 GB if you are using data deduplication. • At least 40 GB for heavily used servers. Using 40 GB or more of memory enhances performance of the Tivoli Storage Manager server database inventory. • If you plan to run multiple instances, each instance requires the memory listed for one server. Multiply the memory for one server by the number of instances planned for the system. • If you plan to use node replication without data deduplication, the system requires 32 GB of memory. Node replication with data deduplication requires a minimum of 64 GB of memory.
Devices	A DVD device that is available for the installation process, if you are installing from DVD media.

Software requirements

Table 4 describes the minimum software requirements that are needed for a server on an HP-UX system.

Table 4. Software requirements

Type of software	Minimum software requirements
Operating system	<p>The HP Itanium system must have operating system 11i v3 (11.31) with the most current maintenance levels installed.</p> <p>11i v3 with:</p> <ul style="list-style-type: none">• PHCO_38658 - libc cumulative patch• PHSS_37202• PHKL_41481• PHKL_42035• PHKL_42335• PHKL_41588• PHSS_41496 <p>The latest available service patches for the operating system must be applied. Older levels without patches do not work with the device drivers that Tivoli Storage Manager uses.</p>
Communication protocol	<p>A communication method that is installed and activated (shared memory is the default).</p>
Device drivers	<p>The Tivoli Storage Manager passthru device driver is used for non-IBM devices. It uses the SCSI passthru interface to communicate with tape drives and tape libraries. The stape device driver is required for tape drives. The schgr device driver is required for tape libraries. The Tivoli Storage Manager device driver package contains device driver tools and ACSLS daemons.</p> <p>For the IBM 3590, 3592, or the Ultrium tape library or drives, the IBM device drivers are required. Install the most current device drivers. You can locate IBM driver packages at Fix Central.</p> <p>Configure the device drivers before you use the server with tape devices.</p>
Gunzip utility	<p>The gunzip utility must be available on your system before you install or upgrade the Version 7 server. Ensure that the gunzip utility is installed and the path to it is set in the PATH environment variable.</p>
Other software	<p>Korn Shell (ksh): You must have the I/O completion ports (IOCP) configured on the operating system.</p> <p>To authenticate Tivoli Storage Manager users with a Lightweight Directory Access Protocol (LDAP) server, you must use one of the following directory servers:</p> <ul style="list-style-type: none">• Microsoft Active Directory (Windows Server 2008 or Windows Server 2012)• IBM Security Directory Server V6.3• IBM Security Directory Server V6.4

Compatibility of the Tivoli Storage Manager server with other DB2 products on the system

You can install other products that deploy and use DB2 products on the same system as the Tivoli Storage Manager Version 7.1.7 server, with some limitations.

To install and use other products that use a DB2 product on the same system as the Tivoli Storage Manager server, ensure that the following criteria are met:

Table 5. Compatibility of the Tivoli Storage Manager server with other DB2 products on the system

Criterion	Instructions
Version level	The other products that use a DB2 product must use DB2 version 9 or later. DB2 products include product encapsulation and segregation support beginning with Version 9. Starting with this version, you can run multiple copies of DB2 products, at different code levels, on the same system. For details, see the information about multiple DB2 copies in the DB2 product information.
User IDs and directories	Ensure that the user IDs, fence user IDs, installation location, other directories, and related information are not shared across DB2 installations. Your specifications must be different from the IDs and locations that you used for the Tivoli Storage Manager server installation and configuration. If you used the dsmicfgx wizard or dsmupgdx wizard to configure Version 7.1.7, or upgrade the server from Version 5.5, these are values that you entered when running the wizard. If you used the manual configuration for Version 7.1.7 or upgrade from Version 5.5 procedures, review the procedures that you used if necessary to recall the values that were used for the server.

Installing the Tivoli Storage Manager server

Table 5. Compatibility of the Tivoli Storage Manager server with other DB2 products on the system (continued)

Criterion	Instructions
Resource allocation	<p>Consider the resources and capability of the system compared to the requirements for both the Tivoli Storage Manager server and the other applications that use the DB2 product. To provide sufficient resources for the other DB2 applications, you might have to change the Tivoli Storage Manager server settings so that the server uses less system memory and resources. Similarly, if the workloads for the other DB2 applications compete with the Tivoli Storage Manager server for processor or memory resources, the performance of the server in handling the expected client workload or other server operations might be adversely affected.</p> <p>To segregate resources and provide more capability for the tuning and allocation of processor, memory, and other system resources for multiple applications, consider using logical partition (LPAR), workload partition (WPAR), or other virtual workstation support. For example, run a DB2 application on its own virtualized system.</p>

IBM Installation Manager

Tivoli Storage Manager uses IBM Installation Manager, which is an installation program that can use remote or local software repositories to install or update many IBM products.

If the required version of IBM Installation Manager is not already installed, it is automatically installed or upgraded when you install Tivoli Storage Manager. It must remain installed on the system so that Tivoli Storage Manager can be updated or uninstalled later as needed.

The following list contains explanations of some terms that are used in IBM Installation Manager:

Offering

An installable unit of a software product.

The Tivoli Storage Manager offering contains all of the media that IBM Installation Manager requires to install Tivoli Storage Manager.

Package

The group of software components that are required to install an offering.

The Tivoli Storage Manager package contains the following components:

- IBM Installation Manager installation program
- Tivoli Storage Manager offering

Package group

A set of packages that share a common parent directory.

Installing the Tivoli Storage Manager server

The default package group for the Tivoli Storage Manager package is IBM Installation Manager.

Repository

A remote or local storage area for data and other application resources.

The Tivoli Storage Manager package is stored in a repository on IBM Fix Central.

Shared resources directory

A directory that contains software files or plug-ins that are shared by packages.

IBM Installation Manager stores installation-related files in the shared resources directory, including files that are used for rolling back to a previous version of Tivoli Storage Manager.

Worksheets for planning details for the server

You can use the worksheets to help you plan the amount and location of storage needed for the Tivoli Storage Manager server. You can also use them to keep track of names and user IDs.

Item	Space required	Number of directories	Location of directories
The database			
Active log			
Archive log			
Optional: Log mirror for the active log			
Optional: Secondary archive log (failover location for archive log)			

Item	Names and user IDs	Location
The <i>instance user ID</i> for the server, which is the ID you use to start and run the Tivoli Storage Manager server		
The <i>home directory</i> for the server, which is the directory that contains the instance user ID		
The database instance name		
The <i>instance directory</i> for the server, which is a directory that contains files specifically for this server instance (the server options file and other server-specific files)		
The server name, use a unique name for each server		

Capacity planning

Capacity planning for Tivoli Storage Manager includes managing resources such as the database, the recovery log and the shared resource area. To maximize resources as part of capacity planning, you must estimate space requirements for the database and the recovery log. The shared resource area must have enough space available for each installation or upgrade.

Estimating space requirements for the database

To estimate space requirements for the database, you can use the maximum number of files that can be in server storage at one time or you can use storage pool capacity.

About this task

Consider using at least 25 GB for the initial database space. Provision file system space appropriately. A database size of 25 GB is adequate for a test environment or a library-manager-only environment. For a production server supporting client workloads, the database size is expected to be larger. If you use random-access disk (DISK) storage pools, more database and log storage space is needed than for sequential-access storage pools.

The maximum size of the Tivoli Storage Manager database is 4 TB.

For information about sizing the database in a production environment that is based on the number of files and on storage pool size, see the following topics.

Estimating database space requirements based on the number of files

If you can estimate the maximum number of files that might be in server storage at a time, you can use that number to estimate space requirements for the database.

About this task

To estimate space requirements that is based on the maximum number of files in server storage, use the following guidelines:

- 600 - 1000 bytes for each stored version of a file, including image backups.

Restriction: The guideline does not include space that is used during data deduplication.

- 100 - 200 bytes for each cached file, copy storage pool file, active-data pool file, and deduplicated file.
- Additional space is required for database optimization to support varying data-access patterns and to support server back-end processing of the data. The amount of extra space is equal to 50% of the estimate for the total number of bytes for file objects.

In the following example for a single client, the calculations are based on the maximum values in the preceding guidelines. The examples do not take into account that you might use file aggregation. In general, when you aggregate small files, it reduces the amount of required database space. File aggregation does not affect space-managed files.

Procedure

1. Calculate the number of file versions. Add each of the following values to obtain the number of file versions:
 - a. Calculate the number of backed-up files. For example, as many as 500,000 client files might be backed up at a time. In this example, storage policies are set to keep up to three copies of backed up files:

$$500,000 \text{ files} * 3 \text{ copies} = 1,500,000 \text{ files}$$
 - b. Calculate the number of archive files. For example, as many as 100,000 client files might be archived copies.
 - c. Calculate the number of space-managed files. For example, as many as 200,000 client files might be migrated from client workstations.

Using 1000 bytes per file, the total amount of database space that is required for the files that belong to the client is 1.8 GB:

$$(1,500,000 + 100,000 + 200,000) * 1000 = 1.8 \text{ GB}$$

2. Calculate the number of cached files, copy storage-pool files, active-data pool files, and deduplicated files:
 - a. Calculate the number of cached copies. For example, caching is enabled in a 5 GB disk storage pool. The high migration threshold of the pool is 90% and the low migration threshold of the pool is 70%. Thus, 20% of the disk pool, or 1 GB, is occupied by cached files.
 If the average file size is about 10 KB, approximately 100,000 files are in cache at any one time:

$$100,000 \text{ files} * 200 \text{ bytes} = 19 \text{ MB}$$
 - b. Calculate the number of copy storage-pool files. All primary storage pools are backed up to the copy storage pool:

$$(1,500,000 + 100,000 + 200,000) * 200 \text{ bytes} = 343 \text{ MB}$$
 - c. Calculate the number of active storage-pool files. All the active client-backup data in primary storage pools is copied to the active-data storage pool. Assume that 500,000 versions of the 1,500,000 backup files in the primary storage pool are active:

$$500,000 * 200 \text{ bytes} = 95 \text{ MB}$$
 - d. Calculate the number of deduplicated files. Assume that a deduplicated storage pool contains 50,000 files:

$$50,000 * 200 \text{ bytes} = 10 \text{ MB}$$

Based on the preceding calculations, about 0.5 GB of extra database space is required for the client's cached files, copy storage-pool files, active-data pool files, and deduplicated files.

3. Calculate the amount of extra space that is required for database optimization. To provide optimal data access and management by the server, extra database space is required. The amount of extra database space is equal to 50% of the total space requirements for file objects.

$$(1.8 + 0.5) * 50\% = 1.2 \text{ GB}$$
4. Calculate the total amount of database space that is required for the client. The total is approximately 3.5 GB:

$$1.8 + 0.5 + 1.2 = 3.5 \text{ GB}$$
5. Calculate the total amount of database space that is required for all clients. If the client that was used in the preceding calculations is typical and you have 500 clients, for example, you can use the following calculation to estimate the total amount of database space that is required for all clients:

$$500 * 3.5 = 1.7 \text{ TB}$$

Results

Tip: In the preceding examples, the results are estimates. The actual size of the database might differ from the estimate because of factors such as the number of directories and the length of the path and file names. Periodically monitor your database and adjust its size as necessary.

What to do next

During normal operations, the Tivoli Storage Manager server might require temporary database space. This space is needed for the following reasons:

- To hold the results of sorting or ordering that are not already being kept and optimized in the database directly. The results are temporarily held in the database for processing.
- To give administrative access to the database through one of the following methods:
 - A DB2 open database connectivity (ODBC) client
 - An Oracle Java database connectivity (JDBC) client
 - Structured Query Language (SQL) to the server from an administrative-client command line

Consider using an extra 50 GB of temporary space for every 500 GB of space for file objects and optimization. See the guidelines in the following table. In the example that is used in the preceding step, a total of 1.7 TB of database space is required for file objects and optimization for 500 clients. Based on that calculation, 200 GB is required for temporary space. The total amount of required database space is 1.9 TB.

Database size	Minimum temporary-space requirement
< 500 GB	50 GB
≥ 500 GB and < 1 TB	100 GB
≥ 1 TB and < 1.5 TB	150 GB
≥ 1.5 and < 2 TB	200 GB
≥ 2 and < 3 TB	250 - 300 GB
≥ 3 and < 4 TB	350 - 400 GB

Estimating database space requirements based on storage pool capacity

To estimate database space requirements based on storage pool capacity, use a ratio of 1 - 5%. For example, if you require 200 TB of storage pool capacity, the size of your database is expected to be 2 - 10 TB. As a general rule, make your database as large as possible to prevent running out of space. If you run out of database space, server operations and client-store operations can fail.

The database manager and temporary space

The Tivoli Storage Manager server database manager manages and allocates system memory and disk space for the database. The amount of database space you require depends on the amount of system memory available and the server workload.

The database manager sorts data in a specific sequence, according to the SQL statement that you issue to request the data. Depending on the workload on the server, and if there is more data than the database manager can manage, the data (that is ordered in sequence) is allocated to temporary disk space. Data is allocated to temporary disk space when there is a large result set. The database manager dynamically manages the memory that is used when data is allocated to temporary disk space.

For example, expiration processing can produce a large result set. If there is not enough system memory on the database to store the result set, some of the data is allocated to temporary disk space. During expiration processing, if a node or file space are selected that are too large to process, the database manager cannot sort the data in memory. The database manager must use temporary space to sort data.

To run database operations, consider adding more database space for the following scenarios:

- The database has a small amount of space and the server operation that requires temporary space uses the remaining free space.
- The file spaces are large, or the file spaces have an assigned policy that creates many file versions.
- The Tivoli Storage Manager server must run with limited memory. The database uses the Tivoli Storage Manager server main memory to run database operations. However, if there is insufficient memory available, the Tivoli Storage Manager server allocates temporary space on disk to the database. For example, if 10G of memory is available and database operations require 12G of memory, the database uses temporary space.
- An out of database space error is displayed when you deploy a Tivoli Storage Manager server. Monitor the server activity log for messages that are related to database space.

Important: Do not change the DB2 software that is installed with the Tivoli Storage Manager installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack, of DB2 software to avoid damage to the database.

Recovery log space requirements

In Tivoli Storage Manager, the term *recovery log* comprises the active log, the archive log, the active log mirror, and the archive failover log. The amount of space that you require for the recovery log depends on various factors, including, for example, the amount of client activity with the server.

Active and archive log space

When you estimate space requirements for active and archive logs, include some extra space for contingencies such as occasional heavy workloads and failovers.

In Tivoli Storage Manager servers V7.1 and later, the active log can be a maximum size of 512 GB. The archive log size is limited to the size of the file system that it is installed on.

Use the following general guidelines when you estimate the size of the active log:

- The suggested starting size for the active log is 16 GB.
- Ensure that the active log is at least large enough for the amount of concurrent activity that the server typically handles. As a precaution, try to anticipate the largest amount of work that the server manages at one time. Provision the active log with extra space that can be used if needed. Consider using 20% of extra space.
- Monitor used and available active log space. Adjust the size of the active log as needed, depending upon factors such as client activity and the level of server operations.
- Ensure that the directory that holds the active log is as large as, or larger than, the size of the active log. A directory that is larger than the active log can accommodate failovers, if they occur.
- Ensure that the file system that contains the active log directory has at least 8 GB of free space for temporary log movement requirements.

The suggested starting size for the archive log is 48 GB.

The archive log directory must be large enough to contain the log files that are generated since the previous full backup. For example, if you perform a full backup of the database every day, the archive log directory must be large enough to hold the log files for all the client activity that occurs during 24 hours. To recover space, the server deletes obsolete archive log files after a full backup of the database. If the archive log directory becomes full and a directory for archive failover logs does not exist, log files remain in the active log directory. This condition can cause the active log directory to fill up and stop the server. When the server restarts, some of the existing active-log space is released.

After the server is installed, you can monitor archive log utilization and the space in the archive log directory. If the space in the archive log directory fills up, it can cause the following problems:

- The server is unable to perform full database backups. Investigate and resolve this problem.
- Other applications write to the archive log directory, exhausting the space that is required by the archive log. Do not share archive log space with other applications including other Tivoli Storage Manager servers. Ensure that each server has a separate storage location that is owned and managed by that specific server.

Example: Estimating active and archive log sizes for basic client-store operations:

Basic client-store operations include backup, archive, and space management. Log space must be sufficient to handle all store transactions that are in progress at one time.

To determine the sizes of the active and archive logs for basic client-store operations, use the following calculation:

$$\begin{array}{l} \text{number of clients} \times \text{files stored during each transaction} \\ \times \text{log space needed for each file} \end{array}$$

This calculation is used in the example in the following table.

Table 6. Basic client-store operations

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3053 bytes	<p>The value of 3053 bytes for each file in a transaction represents the log bytes that are needed when backing up files from a Windows client where the file names are 12 - 120 bytes.</p> <p>This value is based on the results of tests performed under laboratory conditions. The tests consisted of backup-archive clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.</p>
Active log: Suggested size	19.5 GB ¹	<p>Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes.</p> <p>(300 clients x 4096 files stored during each transaction x 3053 bytes for each file) ÷ 1,073,741,824 bytes = 3.5 GB</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>3.5 + 16 = 19.5 GB</p>
Archive log: Suggested size	58.5 GB ¹	<p>Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement.</p> <p>3.5 x 3 = 10.5 GB</p> <p>Increase that amount by the suggested starting size of 48 GB:</p> <p>10.5 + 48 = 58.5 GB</p>

Installing the Tivoli Storage Manager server

Table 6. Basic client-store operations (continued)

Item	Example values	Description
¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log. Monitor your logs and adjust their size if necessary.		

Example: Estimating active and archive log sizes for clients that use multiple sessions:

If the client option RESOURCEUTILIZATION is set to a value that is greater than the default, the concurrent workload for the server increases.

To determine the sizes of the active and archive logs when clients use multiple sessions, use the following calculation:

number of clients x sessions for each client x files stored
during each transaction x log space needed for each file

This calculation is used in the example in the following table.

Table 7. Multiple client sessions

Item	Example values		Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	1000	The number of client nodes that back up, archive, or migrate files every night.
Possible sessions for each client	3	3	The setting of the client option RESOURCEUTILIZATION is larger than the default. Each client session runs a maximum of three sessions in parallel.
Files stored during each transaction	4096	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3053	3053	<p>The value of 3053 bytes for each file in a transaction represents the log bytes needed when backing up files from a Windows client where the file names are 12 - 120 bytes.</p> <p>This value is based on the results of tests performed under laboratory conditions. Tests consisted of clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.</p>

Table 7. Multiple client sessions (continued)

Item	Example values		Description
Active log: Suggested size	26.5 GB ¹	51 GB ¹	<p>The following calculation was used for 300 clients. One GB equals 1,073,741,824 bytes.</p> <p>$(300 \text{ clients} \times 3 \text{ sessions for each client} \times 4096 \text{ files stored during each transaction} \times 3053 \text{ bytes for each file}) \div 1,073,741,824 = 10.5 \text{ GB}$</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>$10.5 + 16 = 26.5 \text{ GB}$</p> <p>The following calculation was used for 1000 clients. One GB equals 1,073,741,824 bytes.</p> <p>$(1000 \text{ clients} \times 3 \text{ sessions for each client} \times 4096 \text{ files store during each transaction} \times 3053 \text{ bytes for each file}) \div 1,073,741,824 = 35 \text{ GB}$</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>$35 + 16 = 51 \text{ GB}$</p>
Archive log: Suggested size	79.5 GB ¹	153 GB ¹	<p>Because of the requirement to be able to store archive logs across three server-database backup cycles, the estimate for the active log is multiplied by 3:</p> <p>$10.5 \times 3 = 31.5 \text{ GB}$</p> <p>$35 \times 3 = 105 \text{ GB}$</p> <p>Increase those amounts by the suggested starting size of 48 GB:</p> <p>$31.5 + 48 = 79.5 \text{ GB}$</p> <p>$105 + 48 = 153 \text{ GB}$</p>
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your active log and adjust its size if necessary.</p>			

Example: Estimating active and archive log sizes for simultaneous write operations:

If client backup operations use storage pools that are configured for simultaneous write, the amount of log space that is required for each file increases.

The log space that is required for each file increases by about 200 bytes for each copy storage pool that is used for a simultaneous write operation. In the example in the following table, data is stored to two copy storage pools in addition to a primary storage pool. The estimated log size increases by 400 bytes for each file. If you use the suggested value of 3053 bytes of log space for each file, the total number of required bytes is 3453.

Installing the Tivoli Storage Manager server

This calculation is used in the example in the following table.

Table 8. Simultaneous write operations

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option TXNGROUPMAX is 4096.
Log space that is required for each file	3453 bytes	<p>3053 bytes plus 200 bytes for each copy storage pool.</p> <p>The value of 3053 bytes for each file in a transaction represents the log bytes that are needed when backing up files from a Windows client where the file names are 12 - 120 bytes.</p> <p>This value is based on the results of tests performed under laboratory conditions. The tests consisted of backup-archive clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.</p>
Active log: Suggested size	20 GB ¹	<p>Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes.</p> <p>(300 clients x 4096 files stored during each transaction x 3453 bytes for each file) ÷ 1,073,741,824 bytes = 4.0 GB</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>4 + 16 = 20 GB</p>
Archive log: Suggested size	60 GB ¹	<p>Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the archive log requirement:</p> <p>4 GB x 3 = 12 GB</p> <p>Increase that amount by the suggested starting size of 48 GB:</p> <p>12 + 48 = 60 GB</p>
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

Example: Estimating active and archive log sizes for basic client store operations and server operations:

Migration of data in server storage, identification processes for data deduplication, reclamation, and expiration might run concurrently with client store operations. Administrative tasks such as administrative commands or SQL queries from administrative clients can also run concurrently with client store operations. Server operations and administrative tasks that run concurrently can increase the active log space that is required.

For example, migration of files from the random-access (DISK) storage pool to a sequential-access disk (FILE) storage pool uses approximately 110 bytes of log space for each file that is migrated. For example, suppose that you have 300 backup-archive clients and each one of them backs up 100,000 files every night. The files are initially stored on DISK and then migrated to a FILE storage pool. To estimate the amount of active log space that is required for the data migration, use the following calculation. The number of clients in the calculation represents the maximum number of client nodes that back up, archive, or migrate files concurrently at any time.

$$300 \text{ clients} \times 100,000 \text{ files for each client} \times 110 \text{ bytes} = 3.1 \text{ GB}$$

Add this value to the estimate for the size of the active log that calculated for basic client store operations.

Example: Estimating active and archive log sizes under conditions of extreme variation:

Problems with running out of active log space can occur if you have many transactions that complete quickly and some transactions that take much longer to complete. A typical case occurs when many workstation or file-server backup sessions are active and a few very large database server-backup sessions are active. If this situation applies to your environment, you might need to increase the size of the active log so that the work completes successfully.

Example: Estimating archive log sizes with full database backups:

The Tivoli Storage Manager server deletes unnecessary files from the archive log only when a full database backup occurs. Consequently, when you estimate the space that is required for the archive log, you must also consider the frequency of full database backups.

For example, if a full database backup occurs once a week, the archive log space must be able to contain the information in the archive log for a full week.

The difference in archive log size for daily and full database backups is shown in the example in the following table.

Table 9. Full database backups

Item	Example values	Description
Maximum number of client nodes that back up, archive, or migrate files concurrently at any time	300	The number of client nodes that back up, archive, or migrate files every night.
Files stored during each transaction	4096	The default value of the server option TXNGROUPMAX is 4096.

Installing the Tivoli Storage Manager server

Table 9. Full database backups (continued)

Item	Example values	Description
Log space that is required for each file	3453 bytes	<p>3053 bytes for each file plus 200 bytes for each copy storage pool.</p> <p>The value of 3053 bytes for each file in a transaction represents the log bytes needed when backing up files from a Windows client where the file names are 12 - 120 bytes.</p> <p>This value is based on the results of tests performed under laboratory conditions. Tests consisted of clients performing backup operations to a random-access disk (DISK) storage pool. DISK pools result in more log use than sequential-access storage pools. Consider a value larger than 3053 bytes if the data being stored has file names that are longer than 12 - 120 bytes.</p>
Active log: Suggested size	20 GB ¹	<p>Use the following calculation to determine the size of the active log. One GB equals 1,073,741,824 bytes.</p> <p>(300 clients x 4096 files per transaction x 3453 bytes per file) ÷ 1,073,741,824 bytes = 4.0 GB</p> <p>Increase that amount by the suggested starting size of 16 GB:</p> <p>4 + 16 = 20 GB</p>
Archive log: Suggested size with a full database backup every day	60 GB ¹	<p>Because of the requirement to be able to store archive logs across three backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement:</p> <p>4 GB x 3 = 12 GB</p> <p>Increase that amount by the suggested starting size of 48 GB:</p> <p>12 + 48 = 60 GB</p>
Archive log: Suggested size with a full database every week	132 GB ¹	<p>Because of the requirement to be able to store archive logs across three server database-backup cycles, multiply the estimate for the active log by 3 to estimate the total archive log requirement. Multiply the result by the number of days between full database backups:</p> <p>(4 GB x 3) x 7 = 84 GB</p> <p>Increase that amount by the suggested starting size of 48 GB:</p> <p>84 + 48 = 132 GB</p>
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that does not use deduplication, 16 GB is the suggested minimum size for an active log. The suggested starting size for an archive log in a production environment that does not use deduplication is 48 GB. If you substitute values from your environment and the results are larger than 16 GB and 48 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>		

Example: Estimating active and archive log sizes for data deduplication operations:

If you deduplicate data, you must consider its effects on space requirements for active and archive logs.

The following factors affect requirements for active and archive log space:

The amount of deduplicated data

The effect of data deduplication on the active log and archive log space depends on the percentage of data that is eligible for deduplication. If the percentage of data that can be deduplicated is relatively high, more log space is required.

The size and number of extents

Approximately 1,500 bytes of active log space are required for each extent that is identified by a duplicate-identification process. For example, if 250,000 extents are identified by a duplicate-identification process, the estimated size of the active log is 358 MB:

250,000 extents identified during each process x 1,500 bytes
for each extent = 358 MB

Consider the following scenario. Three hundred backup-archive clients back up 100,000 files each night. This activity creates a workload of 30,000,000 files. The average number of extents for each file is two. Therefore, the total number of extents is 60,000,000, and the space requirement for the archive log is 84 GB:

60,000,000 extents x 1,500 bytes for each extent = 84 GB

A duplicate-identification process operates on aggregates of files. An aggregate consists of files that are stored in a given transaction, as specified by the TXNGROUPMAX server option. Suppose that the TXNGROUPMAX server option is set to the default of 4096. If the average number of extents for each file is two, the total number of extents in each aggregate is 8192, and the space required for the active log is 12 MB:

8192 extents in each aggregate x 1500 bytes for each extent =
12 MB

The timing and number of the duplicate-identification processes

The timing and number of duplicate-identification processes also affects the size of the active log. Using the 12 MB active-log size that was calculated in the preceding example, the concurrent load on the active log is 120 MB if 10 duplicate-identification processes are running in parallel:

12 MB for each process x 10 processes = 120 MB

File size

Large files that are processed for duplicate identification can also affect the size of the active log. For example, suppose that a backup-archive client backs up an 80 GB, file-system image. This object can have a high number of duplicate extents if, for example, the files included in the file system image were backed up incrementally. For example, assume that a file system image has 1.2 million duplicate extents. The 1.2 million extents in this large file represent a single transaction for a duplicate-identification process. The total space in the active log that is required for this single object is 1.7 GB:

1,200,000 extents x 1,500 bytes for each extent = 1.7 GB

Installing the Tivoli Storage Manager server

If other, smaller duplicate-identification processes occur at the same time as the duplicate-identification process for a single large object, the active log might not have enough space. For example, suppose that a storage pool is enabled for deduplication. The storage pool has a mixture of data, including many relatively small files that range from 10 KB to several hundred KB. The storage pool also has few large objects that have a high percentage of duplicate extents.

To take into account not only space requirements but also the timing and duration of concurrent transactions, increase the estimated size of the active log by a factor of two. For example, suppose that your calculations for space requirements are 25 GB (23.3 GB + 1.7 GB for deduplication of a large object). If deduplication processes are running concurrently, the suggested size of the active log is 50 GB. The suggested size of the archive log is 150 GB.

The examples in the following tables show calculations for active and archive logs. The example in the first table uses an average size of 700 KB for extents. The example in the second table uses an average size of 256 KB. As the examples show, the average deduplicate-extent size of 256 KB indicates a larger estimated size for the active log. To minimize or prevent operational problems for the server, use 256 KB to estimate the size of the active log in your production environment.

Table 10. Average duplicate-extent size of 700 KB

Item	Example values		Description
Size of largest single object to deduplicate	800 GB	4 TB	The granularity of processing for deduplication is at the file level. Therefore, the largest single file to deduplicate represents the largest transaction and a correspondingly large load on the active and archive logs.
Average size of extents	700 KB	700 KB	The deduplication algorithms use a variable block method. Not all deduplicated extents for a given file are the same size, so this calculation assumes an average size for extents.
Extents for a given file	1,198,372 bits	6,135,667 bits	Using the average extent size (700 KB), these calculations represent the total number of extents for a given object. The following calculation was used for an 800 GB object: $(800 \text{ GB} \div 700 \text{ KB}) = 1,198,372 \text{ bits}$ The following calculation was used for a 4 TB object: $(4 \text{ TB} \div 700 \text{ KB}) = 6,135,667 \text{ bits}$
Active log: Suggested size that is required for the deduplication of a single large object during a single duplicate-identification process	1.7 GB	8.6 GB	The estimated active log space that are needed for this transaction.

Table 10. Average duplicate-extent size of 700 KB (continued)

Item	Example values		Description
Active log: Suggested total size	66 GB ¹	79.8 GB ¹	<p>After considering other aspects of the workload on the server in addition to deduplication, multiply the existing estimate by a factor of two. In these examples, the active log space required to deduplicate a single large object is considered along with previous estimates for the required active log size.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(23.3 \text{ GB} + 1.7 \text{ GB}) \times 2 = 50 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $50 + 16 = 66 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(23.3 \text{ GB} + 8.6 \text{ GB}) \times 2 = 63.8 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $63.8 + 16 = 79.8 \text{ GB}$
Archive log: Suggested size	198 GB ¹	239.4 GB ¹	<p>Multiply the estimated size of the active log by a factor of 3.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $50 \text{ GB} \times 3 = 150 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $150 + 48 = 198 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $63.8 \text{ GB} \times 3 = 191.4 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $191.4 + 48 = 239.4 \text{ GB}$
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that uses deduplication, 32 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that uses deduplication is 96 GB. If you substitute values from your environment and the results are larger than 32 GB and 96 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>			

Table 11. Average duplicate-extent size of 256 KB

Item	Example values		Description
Size of largest single object to deduplicate	800 GB	4 TB	The granularity of processing for deduplication is at the file level. Therefore, the largest single file to deduplicate represents the largest transaction and a correspondingly large load on the active and archive logs.

Installing the Tivoli Storage Manager server

Table 11. Average duplicate-extent size of 256 KB (continued)

Item	Example values		Description
Average size of extents	256 KB	256 KB	The deduplication algorithms use a variable block method. Not all deduplicated extents for a given file are the same size, so this calculation assumes an average extent size.
Extents for a given file	3,276,800 bits	16,777,216 bits	<p>Using the average extent size, these calculations represent the total number of extents for a given object.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(800 \text{ GB} \div 256 \text{ KB}) = 3,276,800 \text{ bits}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(4 \text{ TB} \div 256 \text{ KB}) = 16,777,216 \text{ bits}$
Active log: Suggested size that is required for the deduplication of a single large object during a single duplicate-identification process	4.5 GB	23.4 GB	The estimated size of the active log space that is required for this transaction.
Active log: Suggested total size	71.6 GB ¹	109.4 GB ¹	<p>After considering other aspects of the workload on the server in addition to deduplication, multiply the existing estimate by a factor of 2. In these examples, the active log space required to deduplicate a single large object is considered along with previous estimates for the required active log size.</p> <p>The following calculation was used for multiple transactions and an 800 GB object:</p> $(23.3 \text{ GB} + 4.5 \text{ GB}) \times 2 = 55.6 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $55.6 + 16 = 71.6 \text{ GB}$ <p>The following calculation was used for multiple transactions and a 4 TB object:</p> $(23.3 \text{ GB} + 23.4 \text{ GB}) \times 2 = 93.4 \text{ GB}$ <p>Increase that amount by the suggested starting size of 16 GB:</p> $93.4 + 16 = 109.4 \text{ GB}$

Table 11. Average duplicate-extent size of 256 KB (continued)

Item	Example values		Description
Archive log: Suggested size	214.8 GB ¹	328.2 GB ¹	<p>The estimated size of the active log multiplied by a factor of 3.</p> <p>The following calculation was used for an 800 GB object:</p> $55.6 \text{ GB} \times 3 = 166.8 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $166.8 + 48 = 214.8 \text{ GB}$ <p>The following calculation was used for a 4 TB object:</p> $93.4 \text{ GB} \times 3 = 280.2 \text{ GB}$ <p>Increase that amount by the suggested starting size of 48 GB:</p> $280.2 + 48 = 328.2 \text{ GB}$
<p>¹ The example values in this table are used only to illustrate how the sizes for active logs and archive logs are calculated. In a production environment that uses deduplication, 32 GB is the suggested minimum size for an active log. The suggested minimum size for an archive log in a production environment that uses deduplication is 96 GB. If you substitute values from your environment and the results are larger than 32 GB and 96 GB, use your results to size the active log and archive log.</p> <p>Monitor your logs and adjust their size if necessary.</p>			

Active-log mirror space

The active log can be mirrored so that the mirrored copy can be used if the active log files cannot be read. There can be only one active log mirror.

Creating a log mirror is a suggested option. If you increase the size of the active log, the log mirror size is increased automatically. Mirroring the log can affect performance because of the doubled I/O activity that is required to maintain the mirror. The additional space that the log mirror requires is another factor to consider when deciding whether to create a log mirror.

If the mirror log directory becomes full, the server issues error messages to the activity log and to the db2diag.log. Server activity continues.

Archive-failover log space

The archive failover log is used by the server if the archive log directory runs out of space.

Specifying an archive failover log directory can prevent problems that occur if the archive log runs out of space. If both the archive log directory and the drive or file system where the archive failover log directory is located become full, the data remains in the active log directory. This condition can cause the active log to fill up, which causes the server to halt.

Monitoring space utilization for the database and recovery logs

To determine the amount of used and available active log space, you issue the **QUERY LOG** command. To monitor space utilization in the database and recovery logs, you can also check the activity log for messages.

Active log

If the amount of available active log space is too low, the following messages are displayed in the activity log:

ANR4531I: IC_AUTOBACKUP_LOG_USED_SINCE_LAST_BACKUP_TRIGGER

This message is displayed when the active log space exceeds the maximum specified size. The Tivoli Storage Manager server starts a full database backup.

To change the maximum log size, halt the server. Open the `dsmserv.opt` file, and specify a new value for the `ACTIVELOGSIZE` option. When you are finished, restart the server.

ANR0297I: IC_BACKUP_NEEDED_LOG_USED_SINCE_LAST_BACKUP

This message is displayed when the active log space exceeds the maximum specified size. You must back up the database manually.

To change the maximum log size, halt the server. Open the `dsmserv.opt` file, and specify a new value for the `ACTIVELOGSIZE` option. When you are finished, restart the server.

ANR4529I: IC_AUTOBACKUP_LOG_UTILIZATION_TRIGGER

The ratio of used active-log space to available active-log space exceeds the log utilization threshold. If at least one full database backup has occurred, the Tivoli Storage Manager server starts an incremental database backup. Otherwise, the server starts a full database backup.

ANR0295I: IC_BACKUP_NEEDED_LOG_UTILIZATION

The ratio of used active-log space to available active-log space exceeds the log utilization threshold. You must back up the database manually.

Archive log

If the amount of available archive log space is too low, the following message is displayed in the activity log:

ANR0299I: IC_BACKUP_NEEDED_ARCHLOG_USED

The ratio of used archive-log space to available archive-log space exceeds the log utilization threshold. The Tivoli Storage Manager server starts a full automatic database backup.

Database

If the amount of space available for database activities is too low, the following messages are displayed in the activity log:

ANR2992W: IC_LOG_FILE_SYSTEM_UTILIZATION_WARNING_2

The used database space exceeds the threshold for database space utilization. To increase the space for the database, use the **EXTEND DBSPACE** command, the **EXTEND DBSPACE** command, or the **DSMSERV FORMAT** utility with the **DBDIR** parameter.

ANR1546W: FILESYSTEM_DBPATH_LESS_1GB

The available space in the directory where the server database files are located is less than 1 GB.

When a Tivoli Storage Manager server is created with the DSMSESV FORMAT utility or with the configuration wizard, a server database and recovery log are also created. In addition, files are created to hold database information used by the database manager. The path specified in this message indicates the location of the database information used by the database manager. If space is unavailable in the path, the server can no longer function.

You must add space to the file system or make space available on the file system or disk.

Deleting installation rollback files

You can delete certain installation files that were saved during the installation process to free space in the shared resource directory. For example, files that might have been required for a rollback operation are types of files that you can delete.

About this task

To delete the files that are no longer needed, use either the installation graphical wizard or the command line in console mode.

Deleting installation rollback files by using a graphical wizard

You can delete certain installation files that were saved during installation process by using the IBM Installation Manager user interface.

Procedure

1. Open IBM Installation Manager.
2. Click **File > Preferences**.
3. Select **Files for Rollback**.
4. Click **Delete Saved Files** and click **OK**.

Deleting installation rollback files by using the command line

You can delete certain installation files that were saved during the installation process by using the command line.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
eclipse/tools
For example:
/opt/IBM/InstallationManager/eclipse/tools
2. From the tools directory, issue the following command to start an IBM Installation Manager command line:
./imcl -c
3. Enter P to select Preferences.
4. Enter 3 to select Files for Rollback.
5. Enter D to Delete the Files for Rollback.
6. Enter A to Apply Changes and Return to Preferences Menu.
7. Enter C to leave the Preference Menu.

8. Enter X to Exit Installation Manager.

Server naming best practices

Use these descriptions as a reference when you install or upgrade a Tivoli Storage Manager server.

Instance user ID

The instance user ID is used as the basis for other names related to the server instance. The instance user ID is also called the instance owner.

For example: `tsminst1`

The instance user ID is the user ID that must have ownership or read/write access authority to all directories that you create for the database and the recovery log. The standard way to run the server is under the instance user ID. That user ID must also have read/write access to the directories that are used for any **FILE** device classes.

Home directory for the instance user ID

The home directory can be created when creating the instance user ID, by using the option `(-m)` to create a home directory if it does not exist already. Depending on local settings, the home directory might have the form:
`/home/instance_user_ID`

For example: `/home/tsminst1`

The home directory is primarily used to contain the profile for the user ID and for security settings.

Database instance name

The database instance name must be the same as the instance user ID under which you run the server instance.

For example: `tsminst1`

Instance directory

The instance directory is a directory that contains files specifically for a server instance (the server options file and other server-specific files). It can have any name that you want. For easier identification, use a name that ties the directory to the instance name.

You can create the instance directory as a subdirectory of the home directory for the instance user ID. For example: `/home/instance_user_ID/instance_user_ID`

The following example places the instance directory in the home directory for user ID `tsminst1`: `/home/tsminst1/tsminst1`

You can also create the directory in another location, for example: `/tsmserv/tsminst1`

The instance directory stores the following files for the server instance:

- The server options file, `dsmserv.opt`

Installing the Tivoli Storage Manager server

- The server key database file, cert.kdb, and the .arm files (used by clients and other servers to import the Secure Sockets Layer certificates of the server)
- Device configuration file, if the DEVCONFIG server option does not specify a fully qualified name
- Volume history file, if the VOLUMEHISTORY server option does not specify a fully qualified name
- Volumes for **DEVTYPE=FILE** storage pools, if the directory for the device class is not fully specified, or not fully qualified
- User exits
- Trace output (if not fully qualified)

Database name

The database name is always TSMDB1, for every server instance. This name cannot be changed.

Server name

The server name is an internal name for Tivoli Storage Manager, and is used for operations that involve communication among multiple Tivoli Storage Manager servers. Examples include server-to-server communication and library sharing.

Use a unique name for each server. For easy identification from a **QUERY SERVER** command, use a name that reflects the location or purpose of the server.

If you use the wizard, the default name that is suggested is the host name of the system that you are using. You can use a different name that is meaningful in your environment. If you have more than one server on the system and you use the wizard, you can use the default name for only one of the servers. You must enter a unique name for each server.

For example:

PAYROLL
SALES

Directories for database space and recovery log

The directories can be named according to local practices. For easier identification, consider using names that tie the directories to the server instance.

For example, for the archive log:

/tsminst1_archlog

Installation directories

Installation directories for the Tivoli Storage Manager server include the server, DB2, device, language, and other directories. Each one contains several additional directories.

The default directories and their subdirectories are listed here for the server, DB2, devices, and languages:

- **Server directory** (/opt/tivoli/tsm/server/bin), which contains:
 - Server code and licensing
- **Additional server directories:**
 - The command and message help are in the /opt/tivoli/tsm/server/bin directory.
- **DB2 directories**

The DB2 product that is installed as part of the installation of the Tivoli Storage Manager server has the directory structure as documented in DB2 information sources. Protect these directories and files as you do the server directories.

 - /opt/tivoli/tsm/db2
- **Device directories**
 - /opt/tivoli/tsm/devices/bin
- **Language directory**

Language-dependent portions of the program are located here:

/opt/tivoli/tsm/server/bin

You can use US English, German, French, Italian, Spanish, Brazilian Portuguese, Korean, Japanese, traditional Chinese, simplified Chinese, Chinese GBK, Chinese Big5, and Russian.

Chapter 2. Installing the server components

To install the Version 7.1.7 server components, you can use the installation wizard, the command line in console mode, or silent mode.

About this task

Using the Tivoli Storage Manager installation software, you can install the following components:

- Server

Tip: The database (DB2), the Global Security Kit (GSKit) and IBM Java Runtime Environment (JRE) are automatically installed when you select the server component.

- Server languages
- License
- Devices
- Tivoli Storage Manager for Storage Area Networks

Allow approximately 30 - 45 minutes to install a V7.1.7 server, using this guide.

Obtaining the installation package

You can obtain the Tivoli Storage Manager installation package from a DVD or an IBM download site such as Passport Advantage® or IBM Fix Central.

Before you begin

If you plan to download the files, set the system user limit for maximum file size to unlimited to ensure that the files can be downloaded correctly:

1. To query the maximum file size value, issue the following command:

```
ulimit -Hf
```
2. If the system user limit for maximum file size is not set to unlimited, change it to unlimited by following the instructions in the documentation for your operating system.

Procedure

1. Download the appropriate package file from one of the following websites, or you can access the files from the product DVD:
 - For a first time installation or a new release go to Passport Advantage to download the license.
 - For all components except the license, download the server package from Fix Central.
 - For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.
2. If you downloaded the package from an IBM download site, complete the following steps:
 - a. Verify that you have enough space to store the installation files when they are extracted from the product package. See the download document for the space requirements:

Installing the Tivoli Storage Manager server

- Tivoli Storage Manager technote 4042235
 - Tivoli Storage Manager Extended Edition technote 4042236
 - System Storage® Archive Manager technote 4042237
- b. Download the package file to the directory of your choice. The path must contain no more than 128 characters. Be sure to extract the installation files to an empty directory. Do not extract to a directory that contains previously extracted files, or any other files.
 - c. Ensure that executable permission is set for the package. If necessary, change the file permissions by issuing the following command:

```
chmod a+x package_name.bin
```
 - d. Extract the package by issuing the following command:

```
./package_name.bin
```

where *package_name* is the name of the downloaded file, for example:
7.1.x.000-TIV-TSMSRV-HP-UX.bin
3. Ensure that the following command is enabled so that the Tivoli Storage Manager wizards work properly:

```
logins
```

By default, the command is enabled.
 4. Select one of the following methods of installing Tivoli Storage Manager:
 - “Installing Tivoli Storage Manager by using the installation wizard”
 - “Installing Tivoli Storage Manager by using console mode” on page 51
 - “Installing Tivoli Storage Manager in silent mode” on page 52
 5. After you install Tivoli Storage Manager, and before you customize it for your use, go to the IBM Support Portal. Click **Support and downloads** and apply any applicable fixes.

Installing Tivoli Storage Manager by using the installation wizard

You can install the server by using the IBM Installation Manager graphical wizard.

Before you begin

Take the following actions before you start the installation:

- Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.

Procedure

Install Tivoli Storage Manager by using one of the following methods:

Option	Description
Installing the software from a downloaded package:	<ol style="list-style-type: none">1. Change to the directory where you downloaded the package.2. Start the installation wizard by issuing the following command: <pre>./install.sh</pre>

Option	Description
Installing the software from a DVD:	<ol style="list-style-type: none"> 1. Insert the DVD into the DVD drive. Tip: Ensure that the installation files are visible on the DVD drive. 2. Start the installation wizard by issuing the following command from a command line: <code>./install.sh</code>

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM Installation Manager logs directory.
You can view installation log files by clicking **File > View Log** from the Installation Manager tool. To collect these log files, click **Help > Export Data for Problem Analysis** from the Installation Manager tool.
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click **Downloads (fixes and PTFs)** and apply any applicable fixes.
- After you install a new server, review Taking the first steps after you install Tivoli Storage Manager to learn about configuring your server.

Installing Tivoli Storage Manager by using console mode

You can install Tivoli Storage Manager by using the command line in console mode.

Before you begin

Take the following actions before you start the installation:

- Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.

Procedure

Install Tivoli Storage Manager by using one of the following methods:

Option	Description
Installing the software from a downloaded package:	<ol style="list-style-type: none"> 1. Change to the directory where you downloaded the package. 2. Start the installation wizard in console mode by issuing the following command: <code>./install.sh -c</code> <p>Optional: Generate a response file as part of a console mode installation. Complete the console mode installation options, and in the Summary panel, specify G to generate the responses.</p>

Installing the Tivoli Storage Manager server

Option	Description
Installing the software from a DVD:	<ol style="list-style-type: none">1. Insert the DVD into the DVD drive. Tip: Ensure that the installation files are visible on the DVD drive.2. Start the installation wizard in console mode by issuing the following command: <pre>./install.sh -c</pre> Optional: Generate a response file as part of a console mode installation. Complete the console mode installation options, and in the Summary panel, specify G to generate the responses.

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM Installation Manager logs directory, for example:

```
/var/ibm/InstallationManager/logs
```
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click **Downloads (fixes and PTFs)** and apply any applicable fixes.
- After you install a new server, review Taking the first steps after you install Tivoli Storage Manager to learn about configuring your server.

Installing Tivoli Storage Manager in silent mode

You can install or upgrade the server in silent mode. In silent mode, the installation does not send messages to a console but instead stores messages and errors in log files.

Before you begin

To provide data input when you use the silent installation method, you can use a response file. The following sample response files are provided in the input directory where the installation package is extracted:

install_response_sample.xml

Use this file to install the Tivoli Storage Manager components.

update_response_sample.xml

Use this file to upgrade the Tivoli Storage Manager components.

These files contain default values that can help you avoid any unnecessary warnings. To use these files, follow the instructions that are provided in the files.

If you want to customize a response file, you can modify the options that are in the file. For information about response files, see Response files.

Procedure

1. Create a response file. You can modify the sample response file or create your own file.
2. If you install the server and Operations Center in silent mode, create a password for the Operations Center truststore in the response file.

If you are using the `install_response_sample.xml` file, add the password in the following line of the file, where *mypassword* represents the password:

```
<variable name='ssl.password' value='mypassword' />
```

For more information about this password, see Installation checklist

Tip: To upgrade the Operations Center, the truststore password is not required if you are using the `update_response_sample.xml` file.

3. Start the silent installation by issuing the following command from the directory where the installation package is extracted. The value *response_file* represents the response file path and file name:
 - `./install.sh -s -input response_file -acceptLicense`

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM Installation Manager logs directory, for example:
`/var/ibm/InstallationManager/logs`
- After you install the server and components, and before you customize it for your use, go to the IBM Support Portal. Click **Downloads (fixes and PTFs)** and apply any applicable fixes.
- After you install a new server, review Taking the first steps after you install Tivoli Storage Manager to learn about configuring your server.

Installing server language packages

Translations for the server allow the server to display messages and help in languages other than US English. The translations also allow for the use of locale conventions for date, time, and number formatting.

Server language locales

Use either the default language package option or select another language package to display server messages and help.

This language package is automatically installed for the following default language option for Tivoli Storage Manager server messages and help:

- LANGUAGE AMENG

The following system locale must be installed on the system when you use LANGUAGE AMENG:

```
en_US.iso88591
```

For languages or locales other than the default, install the language package that your installation requires.

You can use the languages that are shown:

Table 12. Server languages for HP-UX

Language	LANGUAGE option value
Chinese, Simplified	zh_CN.hp15CN
	zh_CN.utf8

Installing the Tivoli Storage Manager server

Table 12. Server languages for HP-UX (continued)

Language	LANGUAGE option value
Chinese, Traditional	zh_TW.big5
	zh_TW.eucTW
	zh_TW.utf8
English	AMENG (default)
	en_US.utf8
French	fr_FR.iso88591
	fr_FR.utf8
German	de_DE.iso88591
	de_DE.utf8
Italian	it_IT.iso88591
	it_IT.utf8
Japanese	ja_JP.eucJP
	ja_JP.utf8
Korean	ko_KR.eucKR
	ko_KR.utf8
Portuguese, Brazilian	pt_PT.iso88591
	pt_PT.utf8
Russian	ru_RU.iso88595
	ru_RU.utf8
Spanish	es_ES.iso88591
	es_ES.utf8

Configuring a language package

After you configure a language package, messages and help are shown on the server in languages other than US English. Installation packages are provided with Tivoli Storage Manager.

About this task

To set support for a certain locale, complete one of the following tasks:

- Set the LANGUAGE option in the server options file to the name of the locale that you want to use. For example:
To use the `it_IT.iso88591` locale, set the LANGUAGE option to `it_IT.iso88591`.
See “Server language locales” on page 53.
- If you are starting the server in the foreground, set the LC_ALL environment variable to match the value that is set in the server options file. For example, to set the environment variable for Italian, enter the following value:
`export LC_ALL=it_IT`

If the locale is successfully initialized, it formats the date, time, and number for the server. If the locale is not successfully initialized, the server uses the US English message files and the date, time, and number format.

Updating a language package

You can modify or update a language package by using the IBM Installation Manager.

About this task

You can install another language package within the same Tivoli Storage Manager instance.

- Use the **Modify** function of IBM Installation Manager to install another language package.
- Use the **Update** function of IBM Installation Manager to update to newer versions of the language packages.

Tip: In IBM Installation Manager, the term *update* means to discover and install updates and fixes to installed software packages. In this context, *update* and *upgrade* are synonymous.

Installing the Tivoli Storage Manager server

Chapter 3. Taking the first steps after you install Tivoli Storage Manager

After you install Version 7.1.7, prepare for the configuration. Using the configuration wizard is the preferred method of configuring the Tivoli Storage Manager instance.

About this task

1. Update the kernel parameter values.
See “Tuning kernel parameters.”
2. Create the directories and user ID for the server instance. See “Creating the user ID and directories for the server instance” on page 59.
3. Configure a server instance. Select one of the following options:
 - Use the configuration wizard, the preferred method. See “Configuring Tivoli Storage Manager by using the configuration wizard” on page 61.
 - Manually configure the new instance. See “Configuring the server instance manually” on page 62. Complete the following steps during a manual configuration.
 - a. Set up your directories and create the Tivoli Storage Manager instance. See “Creating the server instance” on page 62.
 - b. Create a new server options file by copying the sample file to set up communications between the server and clients. See “Configuring server and client communications” on page 64.
 - c. Issue the **DSMSERV FORMAT** command to format the database. See “Formatting the database and log” on page 67.
 - d. Configure your system for database backup. See “Preparing the database manager for database backup” on page 68.
4. Configure options to control when database reorganization runs. See “Configuring server options for server database maintenance” on page 70.
5. Start the server instance if it is not already started.
See “Starting the server instance” on page 71.
6. Register your license. See “Registering licenses” on page 75.
7. Prepare your system for database backups. See “Specifying a device class in preparation for database backups” on page 75.
8. Monitor the server. See “Monitoring the server” on page 77.

Tuning kernel parameters

For IBM Tivoli Storage Manager and DB2 to install and operate correctly on HP-UX, you must update the kernel configuration parameters.

About this task

If you do not update these parameters, the installation of DB2 and Tivoli Storage Manager might fail. Even if installation is successful, operational problems might occur if you do not set parameter values.

Using the db2osconf utility

Use the **db2osconf** utility to get preferred changes based on the size of your system. The utility checks the settings of the kernel configuration parameters and suggests updates to improve performance.

About this task

The **db2osconf** utility does not change the settings in the `/etc/system` file. You must run the **db2osconf** utility after installing Tivoli Storage Manager. If you update a kernel parameter, you must restart the system.

The suggested updates provided by the **db2osconf** utility are the minimum settings required to run DB2 on your system. In addition to the suggestions made by the **db2osconf** utility, other changes are required to run both Tivoli Storage Manager and DB2.

Procedure

To display a list of options for the **db2osconf** utility, issue the following command:

```
db2osconf -h
```

Results

For details about the **db2osconf** utility, see the DB2 product information. Search for **db2osconf**.

Updating kernel parameters

You must have root authority to update kernel parameters.

About this task

To update kernel parameters on HP-UX, complete the following steps:

Procedure

1. Enter the **sam** command to start the System Administration Manager (SAM) program.
2. Double-click the **Kernel Configuration** icon.
3. Double-click the **Configurable Parameters** icon.
4. Double-click the parameter that you want to change, and type the new value in the Formula/Value field. Click **OK**.
5. Repeat these steps for all of the kernel configuration parameters that you want to change.
6. When you are finished setting all of the kernel configuration parameters, select **Action > Process New Kernel** from the action menu bar.

What to do next

After you change the values for the kernel configuration parameters, the HP-UX operating system automatically restarts.

Suggested minimum values

Ensure that the values for kernel parameters are sufficient to prevent operational problems from occurring when running the Tivoli Storage Manager server.

About this task

The following table contains the suggested minimum values for running Tivoli Storage Manager and DB2.

Parameter	Description	Preferred value
semume	The process-wide maximum number of undo structures per process.	4000
maxdsiz_64bit	The maximum data segment size, in bytes for a running process.	8,589,934,592
msgmnb	The maximum number of bytes on a message queue.	2,097,152
msgmni	The maximum number of message queues.	32,767
msgtql	The maximum number of messages on all message queues.	262,144
shmmni	The maximum number of shared memory segments in the entire system.	Set to at least double the value preferred by the db2osconf utility
shmseg	The maximum number of shared memory segments that one process can attach.	Set to the value the db2osconf utility preferred for shmmni

Creating the user ID and directories for the server instance

Create the user ID for the Tivoli Storage Manager server instance and create the directories that the server instance needs for database and recovery logs.

Before you begin

Review the information about planning space for the server before you complete this task. See “Worksheets for planning details for the server” on page 27.

Procedure

1. Create the user ID that will own the server instance. You use this user ID when you create the server instance in a later step.

Create a user ID and group that will be the owner of the server instance.

- a. The following commands can be run from an administrative user ID that will set up the user and group. Create the user ID and group in the home directory of the user.

Installing the Tivoli Storage Manager server

Restriction: In the user ID, only lowercase letters (a-z), numerals (0-9), and the underscore character (_) can be used. The user ID and group name must comply with the following rules:

- The length must be 8 characters or less.
- The user ID and group name cannot start with *ibm*, *sql*, *sys*, or a numeral.
- The user ID and group name cannot be *user*, *admin*, *guest*, *public*, *local*, or any SQL reserved word.

For example, create user ID `tsminst1` in group `tsmsrvrs`. The following examples show how to create this user ID and group using operating system commands.

```
groupadd tsmsrvrs
useradd -d /home/tsminst1 -m -g tsmsrvrs
        -s /bin/ksh tsminst1
passwd tsminst1
```

Restriction: DB2 does not support direct operating system user authentication through LDAP.

- b. Log off, then log in to your system. Change to the user account that you just created. Use an interactive login program, such as `telnet`, so that you are prompted for the password and can change it if necessary.

2. Create directories that the server requires.

Create empty directories for each item in the table and ensure that the directories are owned by the new user ID you just created. Mount the associated storage to each directory for the active log, archive log, and database directories.

Item	Example commands for creating the directories	Your directories
The <i>instance directory</i> for the server, which is a directory that will contain files specifically for this server instance (the server options file and other server-specific files)	<code>mkdir /tsminst1</code>	
The database directories	<code>mkdir /tsmdb001</code> <code>mkdir /tsmdb002</code> <code>mkdir /tsmdb003</code> <code>mkdir /tsmdb004</code>	
Active log directory	<code>mkdir /tsmlog</code>	
Archive log directory	<code>mkdir /tsmarchlog</code>	
Optional: Directory for the log mirror for the active log	<code>mkdir /tsmlogmirror</code>	
Optional: Secondary archive log directory (failover location for archive log)	<code>mkdir /tsmarchlogfailover</code>	

When a server is initially created by using the **DSMSERV FORMAT** utility or the configuration wizard, a server database and recovery log are created. In addition, files are created to hold database information that is used by the database manager.

3. Log off the new user ID.

Configuring the Tivoli Storage Manager server

After you have installed the server and prepared for the configuration, configure the server instance.

About this task

Configure a Tivoli Storage Manager server instance by selecting one of the following options:

- Use the Tivoli Storage Manager configuration wizard on your local system. See “Configuring Tivoli Storage Manager by using the configuration wizard.”
- Manually configure the new Tivoli Storage Manager instance. See “Configuring the server instance manually” on page 62. Complete the following steps during a manual configuration.
 1. Set up the directories and create the Tivoli Storage Manager instance. See “Creating the server instance” on page 62.
 2. Create a new server options file by copying the sample file in order to set up communications between the Tivoli Storage Manager server and clients. See “Configuring server and client communications” on page 64.
 3. Issue the **DSMSERV FORMAT** command to format the database. See “Formatting the database and log” on page 67.
 4. Configure your system for database backup. See “Preparing the database manager for database backup” on page 68.

Configuring Tivoli Storage Manager by using the configuration wizard

The wizard offers a guided approach to configuring a server. By using the graphical user interface (GUI), you can avoid some configuration steps that are complex when done manually. Start the wizard on the system where you installed the Tivoli Storage Manager server program.

Before you begin

Before you begin to use the configuration wizard, you must complete all preceding steps to prepare for the configuration. These steps include installing Tivoli Storage Manager, creating the database and log directories, and creating the directories and user ID for the server instance.

Procedure

1. Ensure that the following requirements are met:
 - The system where you installed Tivoli Storage Manager must have the X Window System client. You must also be running an X Window System server on your desktop.
 - The system must have the Secure Shell (SSH) protocol enabled. Ensure that the port is set to the default value, 22, and that the port is not blocked by a firewall. You must enable password authentication in the `sshd_config` file in

Installing the Tivoli Storage Manager server

- the `/etc/ssh/` directory. Also, ensure that the SSH daemon service has access rights for connecting to the system by using the `localhost` value.
 - You must be able to log in to Tivoli Storage Manager with the user ID that you created for the server instance, by using the SSH protocol. When you use the wizard, you must provide this user ID and password to access that system.
 - Restart the server before you proceed with the Configuration wizard.
2. Start the local version of the wizard:
 - Open the `dsmicfgx` program in the `/opt/tivoli/tsm/server/bin` directory.
 - This wizard can be only run as a root user.
 - Follow the instructions to complete the configuration. The wizard can be stopped and restarted, but the server is not operational until the entire configuration process is complete.

Configuring the server instance manually

After installing Tivoli Storage Manager Version 7.1.7, you can configure Tivoli Storage Manager manually instead of using the configuration wizard.

Creating the server instance

Create a Tivoli Storage Manager instance by issuing the **db2icrt** command.

About this task

You can have one or more server instances on one workstation.

Important: Before you run the **db2icrt** command, verify the following items:

- The home directory for the user (`/home/tsminst1`) exists. If there is no home directory, you must create it.

The instance directory stores the following core files that are generated by the Tivoli Storage Manager server:

- The server options file, `dsmerv.opt`
- The server key database file, `cert.kdb`, and the `.arm` files (used by clients and other servers to import the Secure Sockets Layer certificates of the server)
- Device configuration file, if the `DEVCONFIG` server option does not specify a fully qualified name
- Volume history file, if the `VOLUMEHISTORY` server option does not specify a fully qualified name
- Volumes for **DEVTYPE=FILE** storage pools, if the directory for the device class is not fully specified, or not fully qualified
- User exits
- Trace output (if not fully qualified)
- A shell configuration file (for example, `.profile`) exists in the home directory. The root user and instance-user ID must have write permission to this file. For more information, see the DB2 product information. Search for Linux and UNIX environment variable settings.

1. Log in using the root user ID and create a Tivoli Storage Manager instance. The name of the instance must be the same name as the user that owns the instance. Use the **db2icrt** command and enter the command on one line:

```
/opt/tivoli/tsm/db2/instance/db2icrt -a server -u  
instance_name instance_name
```


Installing the Tivoli Storage Manager server

For example, if your user ID for this instance is `tsminst1`, use the following command to create the instance. Enter the command on one line.

```
/opt/tivoli/tsm/db2/instance/db2icrt -a server -u  
tsminst1 tsminst1
```

Remember: From this point on, use this new user ID when you configure your Tivoli Storage Manager server. Log out of the root user ID and log in under the new instance-user ID.

2. Change the default directory for the database to be the same as the instance directory for the server. If you have multiple servers, log in under the instance ID for each server. Issue this command:

```
db2 update dbm cfg using dftdbpath instance_directory
```

For example, where `instance_directory` is the instance user ID:

```
db2 update dbm cfg using dftdbpath /tsminst1
```

3. Modify the library path to use the version of the IBM Global Security Kit (GSKit) that is installed with the server. In the following examples, *server_bin_directory* is a subdirectory of the server installation directory. For example, `/opt/tivoli/tsm/server/bin`.

- You must update the following files to set the library path when DB2 or the server are started:

Bash or Korn shell example:

```
instance_users_home_directory/sqllib/userprofile
```

C shell example:

```
instance_users_home_directory/sqllib/usercshrc
```

- Add the following entry to the *instance_users_home_directory*/sqllib/userprofile (Bash or Korn shell) file. Each entry is on one line.

```
LD_LIBRARY_PATH=server_bin_directory/dbbkapi:  
/opt/ibm/gsk8_64/lib64:$LD_LIBRARY_PATH
```

```
export LD_LIBRARY_PATH
```

- Add the following entry to the *instance_users_home_directory*/sqllib/usercshrc (C shell) file, on one line:

```
setenv LD_LIBRARY_PATH server_bin_directory/dbbkapi:  
/opt/ibm/gsk8_64/lib64:$LD_LIBRARY_PATH
```

- Verify the library path settings and that the GSKit is version 8.0.14.43 or later. Issue the following commands:

```
echo $LD_LIBRARY_PATH  
gsk8capicmd_64 -version  
gsk8ver_64
```

If your GSKit version is not 8.0.14.43 or later, you must reinstall the Tivoli Storage Manager server. The reinstallation ensures that the correct GSKit version is available.

4. Create a new server options file. See “Configuring server and client communications” on page 64.

Configuring server and client communications

A default sample server options file, `dsmserv.opt.smp`, is created during Tivoli Storage Manager installation in the `/opt/tivoli/tsm/server/bin` directory. You must set up communications between the server and clients by creating a new server options file. To do so, copy the sample file to the directory for the server instance.

About this task

Ensure that you have a server instance directory, for example `/tsminst1`, and copy the sample file to this directory. Name the new file `dsmserv.opt` and edit the options. Complete this set-up before you initialize the server database. Each sample or default entry in the sample options file is a comment, a line beginning with an asterisk (*). Options are not case-sensitive and one or more blank spaces are allowed between keywords and values.

When editing the options file, follow these guidelines:

- Remove the asterisk at the beginning of the line to activate an option.
- Begin entering the options in any column.
- Enter only one option per line, and the option must be on only one line.
- If you make multiple entries for a keyword, the Tivoli Storage Manager server uses the last entry.

If you change the server options file, you must restart the server for the changes to take effect.

You can specify one or more of the following communication methods:

- TCP/IP Version 4 or Version 6
- Shared memory
- Simple network management protocol (SNMP) DPI
- Secure Sockets Layer (SSL)

Tip: You can authenticate passwords with the LDAP directory server, or authenticate passwords with the Tivoli Storage Manager server. Passwords that are authenticated with the LDAP directory server can provide enhanced system security.

Setting TCP/IP options:

Select from a range of TCP/IP options for the Tivoli Storage Manager server or retain the default.

About this task

The following is an example of a list of TCP/IP options you can use to set up your system.

```
commethod      tcpip
tcpport        1500
tcpwindowsize  0
tcpnodelay     yes
```

Tip: You can use TCP/IP Version 4, Version 6, or both.

TCPPORT

The server TCP/IP port address. The default value is 1500.

TCPWINDOWSIZE

Specifies the size of the TCP/IP buffer that is used when sending or receiving data. The window size that is used in a session is the smaller of the server and client window sizes. Larger window sizes use additional memory but can improve performance.

You can specify an integer from 0 to 2048. To use the default window size for the operating system, specify 0.

TCPNODELAY

Specifies whether or not the server sends small messages or lets TCP/IP buffer the messages. Sending small messages can improve throughput but increases the number of packets sent over the network. Specify YES to send small messages or NO to let TCP/IP buffer them. The default is YES.

TCPADMINPORT

Specifies the port number on which the server TCP/IP communication driver is to wait for requests other than client sessions. The default value is 1500.

SSLTCPPORT

(SSL-only) Specifies the Secure Sockets Layer (SSL) port number on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions for the command-line backup-archive client and the command-line administrative client.

SSLTCPADMINPORT

Specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions for the command-line administrative client.

Setting shared memory options:

You can use shared memory communications between clients and servers on the same system. To use shared memory, TCP/IP Version 4 must be installed on the system.

About this task

Shared memory options is the default communication method.

The following example shows a shared memory setting:

```
commethod    sharedmem
shmport      1510
```

In this example, **SHMPORT** specifies the TCP/IP port address of a server when using shared memory. Use the **SHMPORT** option to specify a different TCP/IP port. The default port address is 1510.

Tip: The IBM Tivoli Storage Manager server supports a maximum of twenty-five concurrent shared memory sessions.

COMMETHOD can be used multiple times in the Tivoli Storage Manager server options file, with a different value each time. For example, the following example is possible:

Installing the Tivoli Storage Manager server

```
commethod tcpip  
commethod sharedmem
```

Ensure that you run the DB2OSCONF utility so that the maximum number of message queues (**MSGMNI**) is correct.

If the server and client are not run under the same user ID, then the server must be root. This prevents shared memory communication errors.

Setting SNMP DPI subagent options:

Tivoli Storage Manager implements a simple network management protocol (SNMP) subagent. You can configure the SNMP subagent to send traps to an SNMP manager, such as NetView®, and to provide support for a Management Information Base (MIB).

About this task

The subagent communicates with the SNMP daemon, which in turn communicates with a management application. The SNMP daemon must support the DPI protocol. Agents are available on AIX. The subagent process is separate from the Tivoli Storage Manager server process, but the subagent gets its information from a server options file. When the SNMP management application is enabled, it can get information and messages from servers.

Use the following SNMP DPI options as an example of a SNMP setting. You must specify the COMMETHOD option.

```
commethod          snmp  
snmpheartbeatinterval 5  
snmpmessagecategory severity
```

Setting Secure Sockets Layer options:

You can add more protection for your data and passwords by using Secure Sockets Layer (SSL).

Before you begin

SSL is the standard technology for creating encrypted sessions between servers and clients. SSL provides a secure channel for servers and clients to communicate over open communication paths. With SSL, the identity of the server is verified through the use of digital certificates.

To ensure better system performance, use SSL only for sessions when it is needed. Consider adding additional processor resources on the Tivoli Storage Manager server to manage the increased requirements.

Formatting the database and log

Use the **DSMSERV FORMAT** utility to initialize a server instance. No other server activity is allowed while you initialize the database and recovery log.

After you set up server communications, you are ready to initialize the database. Ensure that you log in by using the instance user ID. Do not place the directories on file systems that might run out of space. If certain directories (for example, the archive log) become unavailable or full, the server stops. See Capacity planning for more details.

For optimal performance and to facilitate I/O, specify at least two equally sized containers or Logical Unit Numbers (LUNs) for the database. In addition, each active log and archive log should have its own container or LUN.

Setting the exit list handler

Set the **DB2NOEXITLIST** registry variable to ON for each server instance. Log on to the system as the server instance owner and issue this command:

```
db2set -i server_instance_name DB2NOEXITLIST=ON
```

For example:

```
db2set -i tsminst1 DB2NOEXITLIST=ON
```

Initializing a server instance

Use the **DSMSERV FORMAT** utility to initialize a server instance. For example, if the server instance directory is */tsminst1*, issue the following commands:

```
cd /tsminst1
dsmserv format dbdir=/tsmdb001 activelogsiz=32768
activelogdirectory=/activelog archlogdirectory=/archlog
archfailoverlogdirectory=/archfaillog mirrorlogdirectory=/mirrorlog
```


Tip: If you specify multiple directories, ensure that the underlying file systems are of equal size to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.

Tip: If DB2 does not start after you issue the **DSMSERV FORMAT** command, you might need to disable the file system mount option NOSUID. If this option is set on the file system that contains the DB2 instance owner directory, or on any file system that contains the DB2 database, active logs, archive logs, failover logs, or mirrored logs, the option must be disabled to start the system.

After you disable the NOSUID option, remount the file system and then start DB2 by issuing the following command:

```
db2start
```

Related information:

 [DSMSERV FORMAT \(Format the database and log\)](#)

Preparing the database manager for database backup

To back up the data in the database to Tivoli Storage Manager, you must enable the database manager and configure the Tivoli Storage Manager application programming interface (API).

About this task

Starting with Tivoli Storage Manager V7.1, it is no longer necessary to set the API password during a manual configuration of the server. If you set the API password during the manual configuration process, attempts to back up the database might fail.

If you use the configuration wizard to create a Tivoli Storage Manager server instance, you do not have to complete these steps. If you are configuring an instance manually, complete the following steps before you issue either the **BACKUP DB** or the **RESTORE DB** commands.

Attention: If the database is unusable, the entire Tivoli Storage Manager server is unavailable. If a database is lost and cannot be recovered, it might be difficult or impossible to recover data that is managed by that server. Therefore, it is critically important to back up the database.

In the following commands, replace the example values with your actual values. The examples use `tsminst1` for the server instance user ID, `/tsminst1` for the server instance directory, and `/home/tsminst1` as the server instance users home directory.

1. Set the Tivoli Storage Manager API environment-variable configuration for the database instance:
 - a. Log in by using the `tsminst1` user ID.
 - b. When user `tsminst1` is logged in, ensure that the DB2 environment is properly initialized. The DB2 environment is initialized by running the `/home/tsminst1/sqlllib/db2profile` script, which normally runs automatically from the profile of the user ID. Ensure the `.profile` file exists in the instance users home directory, for example, `/home/tsminst1/.profile`. If `.profile` does not run the `db2profile` script, add the following lines:

```
if [ -f /home/tsminst1/sqlllib/db2profile ]; then
    . /home/tsminst1/sqlllib/db2profile
fi
```
 - c. In the `instance_directory/sqlllib/userprofile` file, add the following lines:

```
DSMI_CONFIG=server_instance_directory/tsmdbmgr.opt
DSMI_DIR=server_bin_directory/dbbkapi
DSMI_LOG=server_instance_directory
export DSMI_CONFIG DSMI_DIR DSMI_LOG
```

where:

- *instance_directory* is the home directory of the server instance user.
- *server_instance_directory* is the server instance directory.
- *server_bin_directory* is the server bin directory. The default location is `/opt/tivoli/tsm/server/bin`.

In the `instance_directory/sqlllib/usercshrc` file, add the following lines:

```
setenv DSMI_CONFIG=server_instance_directory/tsmdbmgr.opt
setenv DSMI_DIR=server_bin_directory/dbbkapi
setenv DSMI_LOG=server_instance_directory
```

2. Log off and log in again as `tsminst1`, or issue this command:

. ~/.profile

Tip: Ensure that you enter a space after the initial dot (.) character.

3. Create a file that is named `tsmdbmgr.opt` in the *server_instance* directory, which is in the `/tsminst1` directory in this example, and add the following line:

```
SERVERNAME TSMDBMGR_TSMINST1
```

Remember: The value for `SERVERNAME` must be consistent in the `tsmdbmgr.opt` and `dsm.sys` files.

4. As root user, add the following lines to the Tivoli Storage Manager API `dsm.sys` configuration file. By default, the `dsm.sys` configuration file is in the following default location:

```
server_bin_directory/dbbkapi/dsm.sys
servername TSMDBMGR_TSMINST1
commethod tcpip
tcpserveraddr localhost
tcpport 1500
errorlogname /tsminst1/tsmdbmgr.log
nodename $$_TSMDBMGR_$$
```

where

- *servername* matches the `servername` value in the `tsmdbmgr.opt` file.
 - *commethod* specifies the client API that is used to contact the server for database backup. This value can be `tcpip` or `sharedmem`. For more information about shared memory, see step 5.
 - *tcpserveraddr* specifies the server address that the client API uses to contact the server for database backup. To ensure that the database can be backed up, this value must be `localhost`.
 - *tcpport* specifies the port number that the client API uses to contact the server for database backup. Ensure that you enter the same `tcpport` value that is specified in the `dsmserv.opt` server options file.
 - *errorlogname* specifies the error log where the client API logs errors that are encountered during a database backup. This log is typically in the server instance directory. However, this log can be placed in any location where the instance user ID has write-permission.
 - *nodename* specifies the node name that the client API uses to connect to the server during a database backup. To ensure that the database can be backed up, this value must be `$_TSMDBMGR_`.
5. Optional: Configure the server to back up the database by using shared memory. In this way, you might be able to reduce the processor load and improve throughput. Complete the following steps:

- a. Review the `dsmserv.opt` file. If the following lines are not in the file, add them:

```
commethod sharedmem
shmport port_number
```

where *port_number* specifies the port to be used for shared memory.

- b. In the `dsm.sys` configuration file, locate the following lines:

```
commethod tcpip
tcpserveraddr localhost
tcpport port_number
```

Replace the specified lines with the following lines:

Installing the Tivoli Storage Manager server

```
commmethod sharedmem  
shmport port_number
```

where *port_number* specifies the port to be used for shared memory.

Configuring server options for server database maintenance

To help avoid problems with database growth and server performance, the server automatically monitors its database tables and reorganizes them when needed. Before starting the server for production use, set server options to control when reorganization runs. If you plan to use data deduplication, ensure that the option to run index reorganization is enabled.

About this task

Table and index reorganization requires significant processor resources, active log space, and archive log space. Because database backup takes precedence over reorganization, select the time and duration for reorganization to ensure that the processes do not overlap and reorganization can complete.

If you update these server options while the server is running, you must stop and restart the server before the updated values take effect.

Procedure

1. Modify the server options.

Edit the server options file, `dsmserv.opt`, in the server instance directory. Follow these guidelines when you edit the server options file:

- To enable an option, remove the asterisk at the beginning of the line.
- Enter an option on any line.
- Enter only one option per line. The entire option with its value must be on one line.
- If you have multiple entries for an option in the file, the server uses the last entry.

To view available server options, see the sample file, `dsmserv.opt.smp`, in the `/opt/tivoli/tsm/server/bin` directory.

2. If you plan to use data deduplication, enable the **ALLOWREORGINDEX** server option. Add the following option and value to the server options file:

```
allowreorgindex yes
```

3. Set the **REORGBEGINTIME** and **REORGDURATION** server options to control when reorganization starts and how long it runs. Select a time and duration so that reorganization runs when you expect that the server is least busy. These server options control both table and index reorganization processes.

- a. Set the time for reorganization to start by using the **REORGBEGINTIME** server option. Specify the time by using the 24-hour system. For example, to set the start time for reorganization as 8:30 p.m., specify the following option and value in the server options file:

```
reorgbegintime 20:30
```

- b. Set the interval during which the server can start reorganization. For example, to specify that the server can start reorganization for four hours after the time set by the **REORGBEGINTIME** server option, specify the following option and value in the server options file:

```
reorgduration 4
```


4. If the server was running while you updated the server options file, stop and restart the server.

Related information:

 [ALLOWREORGINDEX](#)

 [ALLOWREORGTABLE](#)

 [REORGBEGINTIME](#)

 [REORGDURATION](#)

Starting the server instance

You can start the server by using the instance user ID, which is the preferred method, or the root user ID.

Before you begin

Ensure that you set access permissions and user limits correctly. For instructions, see “Verifying access rights and user limits.”

About this task

When you start the server by using the instance user ID, you simplify the setup process and avoid potential issues. However, in some cases, it might be necessary to start the server with the root user ID. For example, you might want to use the root user ID to ensure that the server can access specific devices. You can set up the server to start automatically by using either the instance user ID or the root user ID.

Procedure

To start the server, take one of the following actions:

- Start the server by using the instance user ID.
For instructions, see “Starting the server from the instance user ID” on page 73.
- Start the server by using the root user ID.
For instructions about authorizing root user IDs to start the server, see Authorizing root user IDs to start the server (V7.1.1). For instructions about starting the server by using the root user ID, see Starting the server from the root user ID (V7.1.1).
- Start the server automatically.
For instructions, see “Automatically starting servers” on page 73.

Verifying access rights and user limits

Before you start the server, verify access rights and user limits.

About this task

If you do not verify user limits, also known as *ulimits*, you might experience server instability or a failure of the server to respond. You must also verify the system-wide limit for the maximum number of open files. The system-wide limit must be greater than or equal to the user limit.

Installing the Tivoli Storage Manager server

Procedure

1. Verify that the server instance user ID has permissions to start the server.
2. For the server instance that you plan to start, ensure that you have authority to read and write files in the server instance directory. Verify that the `dsmserv.opt` file exists in the server instance directory, and that the file includes parameters for the server instance.
3. If the server is attached to a tape drive, medium changer, or removable media device, and you plan to start the server by using the instance user ID, grant read/write access to the instance user ID for these devices. To set permissions, take one of the following actions:
 - If the system is dedicated to Tivoli Storage Manager and only the Tivoli Storage Manager administrator has access, make the device special file world-writable. On the operating system command line, issue the following command:

```
chmod +w /dev/rmtX
```
 - If the system has multiple users, you can restrict access by making the Tivoli Storage Manager instance user ID the owner of the special device files. On the operating system command line, issue the following command:

```
chmod u+w /dev/rmtX
```
 - If multiple user instances are running on the same system, change the group name, for example TAPEUSERS, and add each Tivoli Storage Manager instance user ID to that group. Then, change the ownership of the device special files to belong to the group TAPEUSERS, and make them group-writable. On the operating system command line, issue the following command:

```
chmod g+w /dev/rmtX
```
4. To prevent server failures during interaction with DB2, tune the kernel parameters.
For instructions about tuning kernel parameters, see “Tuning kernel parameters” on page 57.
5. Verify the following user limits based on the guidelines in the table.

Table 13. User limit (ulimit) values

User limit type	Preferred value	Command to query value
Maximum size of core files created	Unlimited	<code>ulimit -Hc</code>
Maximum size of a data segment for a process	Unlimited	<code>ulimit -Hd</code>
Maximum file size	Unlimited	<code>ulimit -Hf</code>
Maximum number of open files	65536	<code>ulimit -Hn</code>
Maximum amount of processor time in seconds	Unlimited	<code>ulimit -Ht</code>

To modify user limits, follow the instructions in the documentation for your operating system.

Tip: If you plan to start the server automatically by using a script, you can set the user limits in the script.

6. Ensure that the user limit of maximum user processes (the `nproc` setting) is set to the minimum suggested value of 16384.

- a. To verify the current user limit, issue the `ulimit -Hu` command by using the instance user ID. For example:

```
[user@Machine ~]$ ulimit -Hu
16384
```

- b. If the limit of maximum user processes is not set to 16384, set the value to 16384.

Add the following line to the `/etc/security/limits.conf` file:

```
instance_user_id      -      nproc          16384
```

where *instance_user_id* specifies the server instance user ID.

Starting the server from the instance user ID

To start the server from the instance user ID, log in with the instance user ID and issue the appropriate command from the server instance directory.

Before you begin

Ensure that access rights and user limits are set correctly. For instructions, see “Verifying access rights and user limits” on page 71.

Procedure

1. Log in to the system where Tivoli Storage Manager is installed by using the instance user ID for the server.
2. If you do not have a user profile that runs the `db2profile` script, issue the following command:

```
. /home/tsminst1/sqlllib/db2profile
```

Tip: For instructions about updating the user ID login script to run the `db2profile` script automatically, see the DB2 documentation.

3. Start the server by issuing the following command on one line from the server instance directory:

```
usr/bin/dsmserve
```

Tip: The command runs in the foreground so that you can set an administrator ID and connect to the server instance.

For example, if the name of the server instance is `tsminst1` and the server instance directory is `/tsminst1`, you can start the instance by issuing the following commands:

```
cd /tsminst1
. ~/sqlllib/db2profile
/usr/bin/dsmserve
```

Automatically starting servers

You can configure the server to start automatically at system startup. Use the `rc.dsmserve` script, which is provided for this purpose.

Before you begin

Ensure that access rights and user limits are set correctly. For instructions, see “Verifying access rights and user limits” on page 71.

About this task

The **rc.dsmserv** script is in the server installation directory, for example, in the `/opt/tivoli/tsm/server/bin` directory.

Procedure

If you did not use a wizard to configure the server, add an entry to the `/etc/inittab` file for each server that you want to automatically start:

1. Set the run level to the value that corresponds to multiuser mode with networking enabled. Typically, the run level to use is 2, 3, or 5, depending on the operating system and its configuration. Ensure that the run level in the `/etc/inittab` file matches the run level of the operating system. For more information about multiuser mode and run levels, see the documentation for your operating system.
2. On the **rc.dsmserv** command in the `/etc/inittab` file, specify the instance user ID with the `-u` option, and the location of the server instance directory with the `-i` option. If you want to start more than one server instance automatically, add an entry for each server instance. To verify the syntax, see the documentation for your operating system.

Tip: To automatically start a server instance with the root user ID, use the `-U` option.

Example

For example, if the instance owner is `tsminst1` and the server instance directory is `/home/tsminst1/tsminst1`, add the following entry to `/etc/inittab`, on one line:


```
tsm1:2:once:/opt/tivoli/tsm/server/bin/rc.dsmserv -u tsminst1  
-i /home/tsminst1/tsminst1 -q >/dev/console 2>&1
```

In this example, the ID for the process is `tsm1`, and the run level is set to 2.

If you have more than one server instance that you want to run, add an entry for each server instance. For example, if you have instance owner IDs `tsminst1` and `tsminst2`, and instance directories `/home/tsminst1/tsminst1` and `/home/tsminst2/tsminst2`, add the following entries to `/etc/inittab`. Each entry is on one line.

```
tsm1:2:once:/opt/tivoli/tsm/server/bin/rc.dsmserv -u tsminst1  
-i /home/tsminst1/tsminst1 -q >/dev/console 2>&1  
tsm2:2:once:/opt/tivoli/tsm/server/bin/rc.dsmserv -u tsminst2  
-i /home/tsminst2/tsminst2 -q >/dev/console 2>&1
```

Related reference:

 Server startup script: `rc.dsmserv`

Stopping the server

You can stop the server when needed to return control to the operating system. To avoid losing administrative and client node connections, stop the server only after current sessions are completed or canceled.

About this task

To stop the server, issue the following command from the Tivoli Storage Manager command line:

```
halt
```

If you cannot connect to the server with an administrative client and you want to stop the server, you must cancel the process by using the **kill** command with the process ID number (pid). The pid is displayed at initialization.

Important: Before you issue the **kill** command, ensure that you know the correct process ID for the Tivoli Storage Manager server.

The `dsmserv.v6lock` file, in the directory from which the server is running, can be used to identify the process ID of the process to kill. To display the file, enter:

```
cat /instance_dir/dsmserv.v6lock
```

Issue the following command to stop the server:

```
kill -17 dsmserv_pid
```

where *dsmserv_pid* is the process ID number.

Registering licenses

Immediately register any Tivoli Storage Manager licensed functions that you purchase so you do not lose any data after you start server operations, such as backing up your data.

About this task

Use the **REGISTER LICENSE** command for this task.

Example: Register a license

Register the base Tivoli Storage Manager license.

```
register license file=tsmbasic.lic
```

Specifying a device class in preparation for database backups

To prepare the system for automatic and manual database backups, you must specify the device class to be used.

Before you begin

Ensure that you have defined a tape or file device class.

About this task

Complete the following steps to set up your system for database backups.

Installing the Tivoli Storage Manager server

Procedure

1. If you did not use the configuration wizard (dsmicfgx) to configure the server, ensure that you have completed the steps to manually configure the system for database backups.
2. Select the device class to be used for backups of the database. Issue the following command from a IBM Tivoli Storage Manager administrative command line.

```
set dbrecovery device_class_name
```

The device class that you specify is used by the database manager for database backups. If you do not specify a device class with the **SET DBRECOVERY** command, the backup fails.

Example

For example, to specify that the **DBBACK** device class is to be used, issue this command:

```
set dbrecovery dback
```

Running multiple server instances on a single system

You can create more than one server instance on your system. Each server instance has its own instance directory, and database and log directories.

Multiply the memory and other system requirements for one server by the number of instances planned for the system.

The set of files for one instance of the server is stored separately from the files used by another server instance on the same system. Use the steps in “Creating the server instance” on page 62 for each new instance, including creation of the new instance user.

To manage the system memory that is used by each server, use the **DBMEMPERCENT** server option to limit the percentage of system memory. If all servers are equally important, use the same value for each server. If one server is a production server and other servers are test servers, set the value for the production server to a higher value than the test servers.

When you upgrade from Tivoli Storage Manager V6.1, you must upgrade to V6.3 first, and then to V7.1.7. You can upgrade directly from either V6.2 or V6.3 to V7.1. See the upgrade section (Chapter 5, “Upgrading to V7.1.7,” on page 83) for more details. When you upgrade and have multiple servers on your system, you must run the installation wizard only once. The installation wizard collects the database and variables information for all of your original server instances.

If you upgrade from Tivoli Storage Manager V6.2 or V6.3 to V7.1.7 and have multiple servers on your system, all instances that exist in DB2 V9.7 are dropped and recreated in DB2 V10.5. The wizard issues the db2 upgrade *db dbname* command for each database. The database environment variables for each instance on your system are also reconfigured during the upgrade process.

Related tasks:

 Running multiple server instances on a single system (V7.1.1)

Monitoring the server

When you start to use the server in production, monitor the space that is used by the server to ensure that the amount of space is adequate. Adjust the space if needed.

Procedure

1. Monitor the active log to ensure that the size is correct for the workload that is handled by the server instance.

When the server workload reaches its typical expected level, the space that is used by the active log is 80% - 90% of the space that is available to the active log directory. At that point, you might need to increase the amount of space. Whether you must increase the space depends on the types of transactions in the server workload. Transaction characteristics affect how the active log space is used.

The following transaction characteristics can affect the space usage in the active log:

- The number and size of files in backup operations
 - Clients such as file servers that back up large numbers of small files can cause large numbers of transactions that are completed quickly. The transactions might use a large amount of space in the active log, but for a short time.
 - Clients such as a mail server or a database server that back up large amounts of data in few transactions can cause small numbers of transactions that take a long time to complete. The transactions might use a small amount of space in the active log, but for a long time.
- Network connection types
 - Backup operations that occur over fast network connections cause transactions that complete more quickly. The transactions use space in the active log for a shorter time.
 - Backup operations that occur over relatively slower connections cause transactions that take a longer time to complete. The transactions use space in the active log for a longer time.

If the server is handling transactions with a wide variety of characteristics, the space that is used for the active log might increase and decrease significantly over time. For such a server, you might need to ensure that the active log typically has a smaller percentage of its space used. The extra space allows the active log to grow for transactions that take a long time to complete.

2. Monitor the archive log to ensure that space is always available.

Remember: If the archive log becomes full, and the failover archive log becomes full, the active log can become full, and the server stops. The goal is to make enough space available to the archive log so that it never uses all its available space.

You are likely to notice the following pattern:

- a. Initially, the archive log grows rapidly as typical client-backup operations occur.
- b. Database backups occur regularly, either as scheduled or done manually.
- c. After at least two full database backups occur, log pruning occurs automatically. The space that is used by the archive log decreases when the pruning occurs.
- d. Normal client operations continue, and the archive log grows again.

Installing the Tivoli Storage Manager server

- e. Database backups occur regularly, and log pruning occurs as often as full database backups occur.

With this pattern, the archive log grows initially, decreases, and then might grow again. Over time, as normal operations continue, the amount of space that is used by the archive log should reach a relatively constant level.

If the archive log continues to grow, consider taking one or both of these actions:

- Add space to the archive log. You might need to move the archive log to a different file system.
 - Increase the frequency of full database backups, so that log pruning occurs more frequently.
3. If you defined a directory for the failover archive log, determine whether any logs get stored in that directory during normal operations. If the failover log space is being used, consider increasing the size of the archive log. The goal is that the failover archive log is used only under unusual conditions, not in normal operation.

Chapter 4. Installing a Tivoli Storage Manager server fix pack

Tivoli Storage Manager maintenance updates, which are also referred to as fix packs, bring your server up to the current maintenance level.

Before you begin

To install a fix pack or interim fix to the server, install the server at the level on which you want to run it. You do not have to start the server installation at the base release level. For example, if you currently have V7.1.1 installed, you can go directly to the latest fix pack for V7.1. You do not have to start with the V7.1.0 installation if a maintenance update is available.

You must have the Tivoli Storage Manager license package installed. The license package is provided with the purchase of a base release. When you download a fix pack or interim fix from Fix Central, install the server license that is available on the Passport Advantage website. To display messages and help in a language other than US English, install the language package of your choice.

If you upgrade the server to V7.1 or later, and then revert the server to a level that is earlier than V7.1, you must restore the database to a point in time before the upgrade. During the upgrade process, complete the required steps to ensure that the database can be restored: back up the database, the volume history file, the device configuration file, and the server options file. For more information, see Chapter 6, “Reverting from Version 7.1 to the previous V6 server,” on page 91.

Ensure that you retain the installation media from the base release of the installed server. If you installed Tivoli Storage Manager from a DVD, ensure that the DVD is available. If you installed Tivoli Storage Manager from a downloaded package, ensure that the downloaded files are available. If the upgrade fails, and the server license module is uninstalled, the installation media from the server base release are required to reinstall the license.

Visit the IBM Support Portal for the following information:

- A list of the latest maintenance and download fixes. Click **Support and downloads** and apply any applicable fixes.
- Details about obtaining a base license package. Search for **Warranties and licenses**.
- Supported platforms and system requirements. Search for **Tivoli Storage Manager supported operating systems**.

About this task

To install a fix pack or interim fix, complete the following steps.

Attention: Do not alter the DB2 software that is installed with Tivoli Storage Manager installation packages and fix packs. Do not install or upgrade to a different version, release, or fix pack of DB2 software because doing so can damage the database.

Procedure

1. Log in as the root user.

Installing a Tivoli Storage Manager fix pack

2. Obtain the package file for the fix pack or interim fix that you want to install from the IBM Support Portal or from Fix Central.
3. Change to the directory where you placed the executable file and complete the following steps.

Tip: The files are extracted to the current directory. Ensure that the executable file is in the directory where you want the extracted files to be located.

- a. Change file permissions by entering the following command:

```
chmod a+x 7.x.x.x-TIV-TSMALL-platform.bin
```

where *platform* denotes the architecture that Tivoli Storage Manager is to be installed on.

- b. Issue the following command to extract the installation files:

```
./7.x.x.x-TIV-TSMALL-platform.bin
```

4. Back up the database. The preferred method is to use a snapshot backup. A snapshot backup is a full database backup that does not interrupt any scheduled database backups. For example, issue the following Tivoli Storage Manager administrative command:

```
backup db type=dbsnapshot devclass=tapeclass
```

5. Back up the device configuration information. Issue the following Tivoli Storage Manager administrative command:

```
backup devconfig filenames=file_name
```

where *file_name* specifies the name of the file in which to store device configuration information.

6. Save the volume history file to another directory or rename the file. Issue the following Tivoli Storage Manager administrative command:

```
backup volhistory filenames=file_name
```

where *file_name* specifies the name of the file in which to store the volume history information.

7. Save a copy of the server options file, typically named `dsmserv.opt`. The file is in the server instance directory.
8. Halt the server before installing a fix pack or interim fix. Use the **HALT** command.
9. Ensure that extra space is available in the installation directory. The installation of this fix pack might require additional temporary disk space in the installation directory of the server. The amount of additional disk space can be as much as that required for installing a new database as part of a Tivoli Storage Manager installation. The Tivoli Storage Manager installation wizard displays the amount of space that is required for installing the fix pack and the available amount. If the required amount of space is greater than the available amount, the installation stops. If the installation stops, add the required disk space to the file system and restart the installation.
10. Select one of the following ways of installing Tivoli Storage Manager.

Important: After a fix pack is installed, it is not necessary to go through the configuration again. You can stop after completing the installation, fix any errors, then restart your servers.

Install the Tivoli Storage Manager software by using one of the following methods:

Installation wizard

Follow the instructions for your operating system:

“Installing Tivoli Storage Manager by using the installation wizard” on page 50

Tip: After you start the wizard, in the IBM Installation Manager window, click the **Update** icon; do not click the **Install** or **Modify** icon.

Command line in console mode

Follow the instructions for your operating system:

“Installing Tivoli Storage Manager by using console mode” on page 51

Silent mode

Follow the instructions for your operating system:

“Installing Tivoli Storage Manager in silent mode” on page 52

Tip: If you have multiple server instances on your system, run the installation wizard only once. The installation wizard upgrades all server instances.

Results

Correct any errors that are detected during the installation process.

If you installed the server by using the installation wizard, you can view installation logs by using the IBM Installation Manager tool. Click **File > View Log**. To collect log files, from the IBM Installation Manager tool, click **Help > Export Data for Problem Analysis**.

If you installed the server by using console mode or silent mode, you can view error logs in the IBM Installation Manager log directory, for example:

```
/var/ibm/InstallationManager/logs
```

Installing a Tivoli Storage Manager fix pack

Chapter 5. Upgrading to V7.1.7

You can upgrade a V5 or V6 server to V7.1.7.

About this task

To upgrade the server on the same operating system, see the upgrade instructions:

Table 14. Upgrade information


To upgrade from this version	To this version	See this information
V7.1	V7.1 fix pack or interim fix	Chapter 4, "Installing a Tivoli Storage Manager server fix pack," on page 79
V6.2 or V6.3	V7.1.7	"Upgrading from V6.2 or V6.3 to V7.1.7" on page 84
V6.1	V7.1.7	"Upgrading from V6.1 to V7.1.7" on page 90
V5	V7.1.1 or later	Upgrading from Tivoli Storage Manager V5 to V7.1 (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.upgrd.doc/t_srv_upgrd.html)

An upgrade from V6.1.x, V6.2.x, or V6.3.x to V7.1.7 takes approximately 20 - 50 minutes. Your environment might produce different results than that obtained in the labs.

To revert to an earlier version of the server after an upgrade or migration, you must have a full database backup and the installation software for the original server. You must also have key configuration files:

- Volume history file
- Device configuration file
- Server options file
- dsmserv.dsk file (for a server upgrade or migration from V5 to V7.1.7)

Related information:

 [IBM Tivoli Storage Manager Upgrade and Migration Process - Frequently Asked Questions](#)

Upgrading from V6.2 or V6.3 to V7.1.7

You can upgrade the server directly from Tivoli Storage Manager V6.2 or V6.3 to V7.1.7. You do not need to uninstall V6.2 or V6.3.

Before you begin

Ensure that you retain the installation media from the V6.2 or V6.3 server base release that you are upgrading. If you installed the server components from a DVD, ensure that the DVD is available. If you installed the server components from a downloaded package, ensure that the downloaded files are available. If the upgrade fails, and the server license module is uninstalled, the installation media from the server base release are required to reinstall the license.

Procedure

To upgrade the server to V7.1.7, complete the following tasks:

1. “Planning the upgrade”
2. “Preparing the system” on page 85
3. “Installing V7.1.7 and verifying the upgrade” on page 86

Planning the upgrade

Before you upgrade the server from V6.2 or V6.3 to V7.1.7, you must review the relevant planning information, such as system requirements and release notes. Then, select an appropriate day and time to upgrade the system so that you can minimize the impact on production operations.

About this task

In lab tests, the process of upgrading the server from V6.2 or V6.3 to V7.1 took 14 - 45 minutes. The results that you achieve might differ, depending on your hardware and software environment, and the size of the server database.

The following table shows the results that were obtained in lab tests.

Table 15. V6 to V7 upgrade times in lab tests

Version of source system	Version of target system	Operating system	Size of server database	Random access memory	Upgrade time
V6.3.4	V7.1	AIX	17 GB	64 GB	40 minutes
V6.3.4	V7.1	AIX	487 GB	32 GB	45 minutes
V6.3.4	V7.1	AIX	3.8 TB	64 GB	35 minutes
V6.2.5	V7.1	Linux	6.16 GB	16 GB	15 minutes
V6.3.4	V7.1	Linux	30 GB	16 GB	14 minutes
V6.2.5	V7.1	Linux	70 GB	16 GB	24 minutes
V6.3.4	V7.1	Linux	1.4 TB	64 GB	30 minutes
V6.2.5	V7.1	Oracle Solaris	9.43 GB	32 GB	35 minutes
V6.3.4	V7.1	Windows	2.35 TB	64 GB	45 minutes

Procedure

1. Review the hardware and software requirements:

“Minimum system requirements for the Tivoli Storage Manager server” on page 22

For the latest updates related to system requirements, see the Tivoli Storage Manager support website at <http://www.ibm.com/support/docview.wss?uid=swg21243309>.

2. For special instructions or specific information for your operating system, review the release notes (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.common/r_relnotes_srv.html) and readme files (<http://www.ibm.com/support/docview.wss?uid=swg27044931>) for V7.1.7 server components.
3. Select an appropriate day and time to upgrade your system to minimize the impact on production operations. The amount of time that is required to update the system depends on the database size and many other factors. When you start the upgrade process, clients cannot connect to the server until the new software is installed and any required licenses are registered again.

Preparing the system

To prepare the system for the upgrade from V6.2 or V6.3 to V7, you must gather information about each DB2 instance. Then, back up the server database, save key configuration files, cancel sessions, and stop the server.

Procedure

1. Log on to the computer where Tivoli Storage Manager is installed.
Ensure that you are logged on with the instance user ID.
2. Obtain a list of DB2 instances. Issue the following system command:
`/opt/tivoli/tsm/db2/instance/db2ilist`

The output might be similar to the following example:

```
tsminst1
```

Ensure that each instance corresponds to a server that is running on the system.

3. For each DB2 instance, note the default database path, actual database path, database name, database alias, and any DB2 variables that are configured for the instance. Keep the record for future reference. This information is required to restore the V6 database.
4. Connect to the Tivoli Storage Manager server by using an administrative user ID.
5. Back up the Tivoli Storage Manager database by using the **BACKUP DB** command. The preferred method is to create a snapshot backup, which is a full database backup that does not interrupt scheduled database backups. For example, you can create a snapshot backup by issuing the following command:
`backup db type=dbsnapshot devclass=tapeclass`
6. Back up the device configuration information to another directory by issuing the following Tivoli Storage Manager administrative command:
`backup devconfig filenames=file_name`

where *file_name* specifies the name of the file in which to store device configuration information.

Tip: If you decide to restore the V6 database, this file is required.

Upgrading the Tivoli Storage Manager server

7. Back up the volume history file to another directory. Issue the following Tivoli Storage Manager administrative command:

```
backup volhistory filenames=file_name
```

where *file_name* specifies the name of the file in which to store the volume history information.

Tip: If you decide to restore the V6 database, this file is required.

8. Save a copy of the server options file, which is typically named `dsmserv.opt`. The file is in the server instance directory.
9. Prevent activity on the server by disabling new sessions. Issue the following Tivoli Storage Manager administrative commands:

```
disable sessions client  
disable sessions server
```
10. Verify whether any sessions exist, and notify the users that the server will be stopped. To check for existing sessions, issue the following Tivoli Storage Manager administrative command:

```
query session
```
11. Cancel sessions by issuing the following Tivoli Storage Manager administrative command:

```
cancel session all
```

This command cancels all sessions except for your current session.

12. Stop the server by issuing the following Tivoli Storage Manager administrative command:

```
halt
```
13. Verify that the server is shut down and no processes are running. Issue the following command:

```
ps -ef | grep dsmserv
```
14. In the server instance directory of your installation, locate the `NODELOCK` file and move it to another directory, where you are saving configuration files. The `NODELOCK` file contains the previous licensing information for your installation. This licensing information is replaced when the upgrade is complete.

Installing V7.1.7 and verifying the upgrade

To complete the process of upgrading the server from V6.2 or V6.3 to V7.1.7, you must install the V7.1.7 server. Then, verify that the upgrade was successful by starting the server instance.

Before you begin

You must be logged on to the system by using the root user ID.

You can obtain the installation package from the product DVD or from an IBM download site.

If you plan to download the files, set the system user limit for maximum file size to unlimited to ensure that the files can be downloaded correctly.

1. To query the maximum file size value, issue the following command:

```
ulimit -Hf
```


2. If the system user limit for maximum file size is not set to unlimited, change it to unlimited by following the instructions in the documentation for your operating system.

About this task

By using the Tivoli Storage Manager installation software, you can install the following components:

- Server

Tip: The database (DB2), the Global Security Kit (GSKit) and IBM Java Runtime Environment (JRE) are automatically installed when you select the server component.

- Server languages
- License
- Devices
- Tivoli Storage Manager for Storage Area Networks
- Operations Center

Procedure

1. If you are obtaining the package from an IBM download site, download the appropriate package file from one of the following websites:
 - For a first time installation or a new release go to Passport Advantage to download the license.
 - For all components except the license, download the server package from Fix Central.
 - For the latest information, updates, and maintenance fixes, go to the IBM Support Portal.
2. If you are downloading the package from one of the download sites, complete the following steps:

- a. Verify that you have enough space to store the installation files when they are extracted from the product package. For space requirements, see the download document for your product.
 - Tivoli Storage Manager technote 4042235
 - Tivoli Storage Manager Extended Edition technote 4042236
 - System Storage Archive Manager technote 4042237
- b. Download the package file to the directory of your choice. The path must contain no more than 128 characters. Be sure to extract the installation files to an empty directory. Do not extract to a directory that contains previously extracted files, or any other files. Also, ensure that you have executable permission for the package file.
- c. If necessary, change the file permissions by issuing the following command:

```
chmod a+x package_name.bin
```

where *package_name* is like the following example:

```
7.1.x.000-TIV-TSMSRV-HP-UX.bin
```

In the examples, 7.1.x.000 represents the product release level.

Upgrading the Tivoli Storage Manager server

- d. Extract the installation files by issuing the following command:
`./package_name.bin`

The package is large. Therefore, the extraction takes some time.

3. To ensure that the Tivoli Storage Manager wizards work correctly, verify that the following command is enabled:

`logins`

By default, the command is enabled.

4. Install the Tivoli Storage Manager software by using one of the following methods. During the installation process, you must install the Tivoli Storage Manager license.

Tip: If you have multiple server instances on your system, install the Tivoli Storage Manager software only once to upgrade all server instances.

Installation wizard

To install the server by using the graphical wizard of IBM Installation Manager, follow the instructions in “Installing Tivoli Storage Manager by using the installation wizard” on page 50.

Ensure that your system meets the prerequisites for using the installation wizard. Then, complete the installation procedure. In the IBM Installation Manager window, click the **Install** icon; do not click the **Update** or **Modify** icon.

Command line in console mode

To install the server by using the command line in console mode, follow the instructions in “Installing Tivoli Storage Manager by using console mode” on page 51.

Review the information about installing the server in console mode and then complete the installation procedure.

Silent mode

To install the server by using silent mode, follow the instructions in “Installing Tivoli Storage Manager in silent mode” on page 52.

Review the information about installing the server in silent mode and then complete the installation procedure.

After you install the software, you do not have to reconfigure the system.

5. Correct any errors that are detected during the installation process.
If you installed the server by using the installation wizard, you can view installation logs by using the IBM Installation Manager tool. Click **File > View Log**. To collect log files, from the IBM Installation Manager tool, click **Help > Export Data for Problem Analysis**.
If you installed the server by using console mode or silent mode, you can view error logs in the IBM Installation Manager log directory, for example:
`/var/ibm/InstallationManager/logs`
6. Obtain any applicable fixes by going to the IBM Support Portal. Click **Downloads (fixes and PTFs)** and apply any applicable fixes.
7. Verify that the upgrade was successful:
 - a. Start the server instance.

For instructions, see “Starting the server instance” on page 71.

- b. Monitor the messages that the server issues as it starts. Watch for error and warning messages, and resolve any issues.
- c. Verify that you can connect to the server by using the administrative client. To start an administrative client session, issue the following Tivoli Storage Manager administrative command:

```
dsmadmcc
```

- d. To obtain information about the upgraded system, run **QUERY** commands. For example, to obtain consolidated information about the system, issue the following Tivoli Storage Manager administrative command:

```
query system
```

To obtain information about the database, issue the following Tivoli Storage Manager administrative command:

```
query db format=detailed
```

- 8. Register the licenses for the Tivoli Storage Manager server components that are installed on your system by issuing the **REGISTER LICENSE** command:

```
register license file=installation_directory/server/bin/component_name.lic
```

where *installation_directory* specifies the directory in which you installed the component, and *component_name* specifies the abbreviation for the component.

For example, if you installed the server in the default directory, /opt/tivoli/tsm, register the license by issuing the following command:

```
register license file=/opt/tivoli/tsm/server/bin/tsmbasic.lic
```

For example, if you installed Tivoli Storage Manager Extended Edition in the /opt/tivoli/tsm directory, issue the following command:

```
register license file=/opt/tivoli/tsm/server/bin/tsmee.lic
```

For example, if you installed System Storage Archive Manager in the /opt/tivoli/tsm directory, issue the following command:

```
register license file=/opt/tivoli/tsm/server/bin/dataret.lic
```

Restriction: You cannot use the Tivoli Storage Manager server to register licenses for IBM Tivoli Storage Manager for Mail, Tivoli Storage Manager for Databases, IBM Tivoli Storage Manager for Enterprise Resource Planning, and Tivoli Storage Manager for Space Management. The **REGISTER LICENSE** command does not apply to these licenses. The licensing for these products is done by Tivoli Storage Manager clients.

- 9. Optional: To install an additional language package, use the modify function of the IBM Installation Manager.
- 10. Optional: To upgrade to a newer version of a language package, use the update function of the IBM Installation Manager.

What to do next

You can authenticate passwords with the LDAP directory server, or authenticate passwords with the Tivoli Storage Manager server. Passwords that are authenticated with the LDAP directory server can provide enhanced system security.

Upgrading from V6.1 to V7.1.7

You can upgrade the Tivoli Storage Manager server from V6.1 to V6.3, and then upgrade the server to V7.1.7. It is not possible to upgrade the server directly from V6.1 to V7.1.7.

Before you begin

Ensure that you retain the installation media from the base release of the V6.1 and V6.3 servers. If you obtained the Tivoli Storage Manager software from a DVD, ensure that the DVD is available. If you obtained the Tivoli Storage Manager software from a downloaded package, ensure that the downloaded files are available. If the upgrade fails, and the server license module is uninstalled, the installation media from the server base release are required to reinstall the license.

Procedure

1. Upgrade the server from V6.1 to V6.3, as described in Upgrading from Tivoli Storage Manager Version 6.1 to 6.3 or later.
2. Upgrade the server from V6.3 to V7.1.7, as described in “Upgrading from V6.2 or V6.3 to V7.1.7” on page 84.

Chapter 6. Reverting from Version 7.1 to the previous V6 server

If you must revert to the previous version of the server after an upgrade, you must have a full database backup from your original version. You must also have the server installation media for your original version and key configuration files. Carefully follow the preparation steps before you upgrade the server. By doing so, it might be possible to revert to the previous version of the Tivoli Storage Manager server with minimal loss of data.

Before you begin

You must have the following items from the earlier version of the server:

- Server database backup
- Volume history file
- Device configuration file
- Server options file

About this task

Use the same instructions whether you are reverting within releases or to an earlier release, for example, from 6.2.2 to 6.2.0 or from 6.2.2 to 6.1.2. The older version must match the version that you used before the upgrade to 7.1.7.

Attention: Specify the **REUSEDELAY** parameter to help prevent backup-archive client data loss when you revert the server to a previous version.

Steps for reverting to the previous server version

About this task

Complete the following steps on the system that has the V7.1 server.

Procedure

1. Halt the server to shut down all server operations by using the **HALT** command.
2. Remove the database from the database manager, then delete the database and recovery log directories.
 - a. Manually remove the database. One way to remove it is by issuing this command:

```
dsmserv removedb tsmdb1
```
 - b. If you must reuse the space that is occupied by the database and recovery log directories, you can now delete these directories.
3. Use the uninstallation program to uninstall the V7.1.7 server. Uninstallation removes the server and the database manager, with their directories. For details, see Chapter 8, "Uninstalling Tivoli Storage Manager," on page 99.
4. Reinstall the version of the server program that you were using before the upgrade to V7.1. This version must match the version that your server was running when you created the database backup that you restore in a later step. For example, the server was at version 6.2.2.0 before the upgrade, and

Reverting to a previous Version 6 server version

you intend to use the database backup that was in use on this server. You must install the 6.2.2.0 fix pack to be able to restore the database backup.

5. Copy the following files to the instance directory.
 - Device configuration file
 - Volume history file
 - The server options file (typically `dsmserv.opt`)
6. Configure the new server database by using the configuration wizard. To start the wizard, issue the following command:

```
. /dsmicfgx
```
7. Ensure that no servers are running in the background.
8. Restore the database to a point in time before the upgrade.
9. If you enabled data deduplication for any FILE-type storage pools that existed before the upgrade, or if you moved data that existed before the upgrade into new storage pools while using the V7.1 server, you must complete additional recovery steps. For more details, see “Additional recovery steps if you created new storage pools or enabled data deduplication.”
10. If the **REUSEDELAY** parameter setting on storage pools is less than the age of the database that you restored, restore volumes on any sequential-access storage pools that were reclaimed after that database backup. Use the **RESTORE VOLUME** command.

If you do not have a backup of a storage pool, audit the reclaimed volumes by using the **AUDIT VOLUME** command, with the **FIX=YES** parameter to resolve inconsistencies. For example:

```
audit volume volume_name fix=yes
```
11. If client backup or archive operations were completed using the V7.1 server, audit the storage pool volumes on which the data was stored.

Additional recovery steps if you created new storage pools or enabled data deduplication

If you created new storage pools, turned on data deduplication for any FILE-type storage pools, or did both while your server was running as a V7.1.7 server, you must complete more steps to return to the previous server version.

Before you begin

To complete this task, you must have a complete backup of the storage pool that was created before the upgrade to V7.1.7.

About this task

Use this information if you did either or both of the following actions while your server was running as a V7.1.7 server:

- You enabled the data deduplication function for any storage pools that existed before the upgrade to V7.1.7 program. Data deduplication applies only to storage pools that use a FILE device type.
- You created new primary storage pools after the upgrade *and* moved data that was stored in other storage pools into the new storage pools.

Complete these steps after the server is again restored to V6.

Procedure

- For each storage pool for which you enabled the data deduplication function, restore the entire storage pool by using the **RESTORE STGPPOOL** command.
- For storage pools that you created after the upgrade, determine what action to take. Data that was moved from existing V6 storage pools into the new storage pools might be lost because the new storage pools no longer exist in your restored V6 server. Possible recovery depends on the type of storage pool:
 - If data was moved from V6 DISK-type storage pools into a new storage pool, space that was occupied by the data that was moved was probably reused. Therefore, you must restore the original V6 storage pools by using the storage pool backups that were created before the upgrade to V7.1.
If *no* data was moved from V6 DISK-type storage pools into a new storage pool, then audit the storage pool volumes in these DISK-type storage pools.
 - If data was moved from V6 sequential-access storage pools into a new storage pool, that data might still exist and be usable in storage pool volumes on the restored V6 server. The data might be usable if the **REUSEDELAY** parameter for the storage pool was set to a value that prevented reclamation while the server was running as a V7.1.7 server. If any volumes were reclaimed while the server was running as a V7.1.7 server, restore those volumes from storage pool backups that were created before the upgrade to V7.1.7.

Reverting to a previous Version 6 server version

Chapter 7. Reference: DB2 commands for Tivoli Storage Manager server databases

Use this list as reference when you are directed to issue DB2 commands by IBM support.

Purpose

After using the wizards to install and configure Tivoli Storage Manager, you seldom need to issue DB2 commands. A limited set of DB2 commands that you might use or be asked to issue are listed in Table 16. This list is supplemental material only and is not a comprehensive list. There is no implication that a Tivoli Storage Manager administrator will use it on a daily or ongoing basis. Samples of some commands are provided. Details of output are not listed.

For a full explanation of the commands described here and of their syntax, see the DB2 product information.

Table 16. DB2 commands

Command	Description	Example
db2icrt	Creates DB2 instances in the home directory of the instance owner. Tip: The Tivoli Storage Manager configuration wizard creates the instance used by the server and database. After a server is installed and configured through the configuration wizard, the db2icrt command is generally not used. This utility is in the DB2DIR/instance directory, where DB2DIR represents the installation location where the current version of the DB2 database system is installed.	Manually create a Tivoli Storage Manager instance. Enter the command on one line: <code>/opt/tivoli/tsm/db2/instance/ db2icrt -a server -u instance_name instance_name</code>
db2osconf	Checks the settings of the kernel configuration parameters and suggests updates to improve performance. The suggested updates provided by the db2osconf utility are the minimum settings required to run DB2 on your system.	Retrieve a list of currently supported options: <code>db2osconf</code>
db2set	Displays DB2 variables.	List DB2 variables: <code>db2set</code>
CATALOG DATABASE	Stores database location information in the system database directory. The database can be located either on the local workstation or on a remote database partition server. The server configuration wizard takes care of any catalog needed for using the server database. Run this command manually, after a server is configured and running, only if something in the environment changes or is damaged.	Catalog the database: <code>db2 catalog database tsmdb1</code>
CONNECT TO DATABASE	Connects to a specified database for command-line interface (CLI) use.	Connect to the Tivoli Storage Manager database from a DB2 CLI: <code>db2 connect to tsmdb1</code>

Reference: DB2 commands for Tivoli Storage Manager server databases

Table 16. DB2 commands (continued)

Command	Description	Example
GET DATABASE CONFIGURATION	Returns the values of individual entries in a specific database configuration file. Important: This command and parameters are set and managed directly by DB2. They are listed here for informational purposes and a means to view the existing settings. Changing these settings might be advised by IBM support or through service bulletins such as APARs or Technical Guidance documents (technotes). Do not change these settings manually. Change them only at the direction of IBM and only through the use of Tivoli Storage Manager server commands or procedures.	Show the configuration information for a database alias: db2 get db cfg for tsmdb1 Retrieve information in order to verify settings such as database configuration, log mode, and maintenance. db2 get db config for tsmdb1 show detail
GET DATABASE MANAGER CONFIGURATION	Returns the values of individual entries in a specific database configuration file. Important: This command and parameters are set and managed directly by DB2. They are listed here for informational purposes and a means to view the existing settings. Changing these settings might be advised by IBM support or through service bulletins such as APARs or Technical Guidance documents (technotes). Do not change these settings manually. Change them only at the direction of IBM and only through the use of Tivoli Storage Manager server commands or procedures.	Retrieve configuration information for the database manager: db2 get dbm cfg
GET HEALTH SNAPSHOT	Retrieves the health status information for the database manager and its databases. The information returned represents a snapshot of the health state at the time the command was issued. Tivoli Storage Manager monitors the state of the database using the health snapshot and other mechanisms that are provided by DB2. There might be cases where the health snapshot or other DB2 documentation indicates that an item or database resource might be in an alert state. Such a case indicates that action must be considered to remedy the situation. Tivoli Storage Manager monitors the condition and responds appropriately. Not all declared alerts by the DB2 database are acted on.	Receive a report on DB2 health monitor indicators: db2 get health snapshot for database on tsmdb1
GRANT (Database Authorities)	Grants authorities that apply to the entire database rather than privileges that apply to specific objects within the database.	Grant access to the user ID itmuser: db2 GRANT CONNECT ON DATABASE TO USER itmuser db2 GRANT CREATETAB ON DATABASE TO USER itmuser

Table 16. DB2 commands (continued)

Command	Description	Example
RUNSTATS	<p>Updates statistics about the characteristics of a table and associated indexes or statistical views. These characteristics include number of records, number of pages, and average record length.</p> <p>To see a table, issue this utility after updating or reorganizing the table.</p> <p>A view must be enabled for optimization before its statistics can be used to optimize a query. A view that is enabled for optimization is known as a statistical view. Use the DB2 ALTER VIEW statement to enable a view for optimization. Issue the RUNSTATS utility when changes to underlying tables substantially affect the rows returned by the view.</p> <p>Tip: The server configures DB2 to run the RUNSTATS command as needed.</p>	<p>Update statistics on a single table.</p> <pre>db2 runstats on table SCHEMA_NAME.TABLE_NAME with distribution and sampled detailed indexes all</pre>
SET SCHEMA	<p>Changes the value of the CURRENT SCHEMA special register, in preparation for issuing SQL commands directly through the DB2 CLI.</p> <p>Tip: A special register is a storage area that is defined for an application process by the database manager. It is used to store information that can be referenced in SQL statements.</p>	<p>Set the schema for Tivoli Storage Manager:</p> <pre>db2 set schema tsmdb1</pre>
START DATABASE MANAGER	<p>Starts the current database manager instance background processes. The server starts and stops the instance and database whenever the server starts and halts.</p> <p>Important: Allow the server to manage the starting and stopping of the instance and database unless otherwise directed by IBM support.</p>	<p>Start the database manager:</p> <pre>db2start</pre>
STOP DATABASE MANAGER	<p>Stops the current database manager instance. Unless explicitly stopped, the database manager continues to be active. This command does not stop the database manager instance if any applications are connected to databases. If there are no database connections, but there are instance attachments, the command forces the instance attachments to stop first. Then, it stops the database manager. This command also deactivates any outstanding database activations before stopping the database manager.</p> <p>This command is not valid on a client.</p> <p>The server starts and stops the instance and database whenever the server starts and halts.</p> <p>Important: Allow the server to manage the starting and stopping of the instance and database unless otherwise directed by IBM support.</p>	<p>Stop the database manager:</p> <pre>db2 stop dbm</pre>

Chapter 8. Uninstalling Tivoli Storage Manager

You can use the following procedures to uninstall Tivoli Storage Manager. Before you remove Tivoli Storage Manager, ensure that you do not lose your backup and archive data.

Before you begin

Complete the following steps before you uninstall Tivoli Storage Manager:

- Complete a full database backup.
- Save a copy of the volume history and device configuration files.
- Store the output volumes in a safe location.

About this task

You can uninstall Tivoli Storage Manager by using any of the following methods: a graphical wizard, the command line in console mode, or silent mode.

“Uninstalling Tivoli Storage Manager by using a graphical wizard”

“Uninstalling Tivoli Storage Manager in console mode” on page 100

“Uninstalling Tivoli Storage Manager in silent mode” on page 100

What to do next

See Chapter 2, “Installing the server components,” on page 49 for installation steps to reinstall the Tivoli Storage Manager components.

Uninstalling Tivoli Storage Manager by using a graphical wizard

You can uninstall Tivoli Storage Manager by using the IBM Installation Manager installation wizard.

Procedure

1. Start the Installation Manager.
In the directory where the Installation Manager is installed, go to the eclipse subdirectory (for example, /opt/IBM/InstallationManager/eclipse), and issue the following command:

```
./IBMIM
```
2. Click **Uninstall**.
3. Select **Tivoli Storage Manager server**, and click **Next**.
4. Click **Uninstall**.
5. Click **Finish**.

Uninstalling Tivoli Storage Manager in console mode

To uninstall Tivoli Storage Manager by using the command line, you must run the uninstallation program of IBM Installation Manager from the command line with the parameter for console mode.

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
eclipse/tools
For example:
/opt/IBM/InstallationManager/eclipse/tools
2. From the tools directory, issue the following command:
./imcl -c
3. To uninstall, enter 5.
4. Choose to uninstall from the IBM Tivoli Storage Manager package group.
5. Enter N for Next.
6. Choose to uninstall the IBM Tivoli Storage Manager server package.
7. Enter N for Next.
8. Enter U for Uninstall.
9. Enter F for Finish.

Uninstalling Tivoli Storage Manager in silent mode

To uninstall Tivoli Storage Manager in silent mode, you must run the uninstallation program of IBM Installation Manager from the command line with the parameters for silent mode.

Before you begin

You can use a response file to provide data input to silently uninstall the Tivoli Storage Manager server components. Tivoli Storage Manager includes a sample response file, `uninstall_response_sample.xml`, in the input directory where the installation package is extracted. This file contains default values to help you avoid any unnecessary warnings.

If you want to uninstall all Tivoli Storage Manager components, leave `modify="false"` set for each component in the response file. If you do not want to uninstall a component, set the value to `modify="true"`.

If you want to customize the response file, you can modify the options that are in the file. For information about response files, see [Response files](#).

Procedure

1. In the directory where IBM Installation Manager is installed, go to the following subdirectory:
eclipse/tools
For example:
/opt/IBM/InstallationManager/eclipse/tools
2. From the tools directory, issue the following command, where *response_file* represents the response file path, including the file name:

```
./imcl -input response_file -silent
```

The following command is an example:

```
./imcl -input /tmp/input/uninstall_response.xml -silent
```

Uninstalling and reinstalling Tivoli Storage Manager

If you plan to manually reinstall Tivoli Storage Manager instead of using the wizard, there are a number of steps to take to preserve your server instance names and database directories. During an uninstallation, any server instances previously set up are removed, but the database catalogs for those instances still exist.

About this task

If you are using the wizard to upgrade from Tivoli Storage Manager Version 6.2 or Version 6.3, it is not necessary to complete these steps. The wizard completes them automatically. To manually uninstall and reinstall Tivoli Storage Manager, complete the following steps:

1. Make a list of your current server instances before proceeding to the uninstallation. Run the following command:

```
/opt/tivoli/tsm/db2/instance/db2ilist
```

2. Run the following commands for every server instance:

```
db2 attach to instance_name
db2 get dbm cfg show detail
db2 detach
```

Keep a record of the database path for each instance.

3. Uninstall Tivoli Storage Manager. See Chapter 8, “Uninstalling Tivoli Storage Manager,” on page 99.
4. When you uninstall any version of Tivoli Storage Manager 6.1 or later, including a fix pack, an instance file is created. The instance file is created to help reinstall Tivoli Storage Manager. Check this file and use the information when you are prompted for the instance credentials when reinstalling. In silent installation mode, you provide these credentials using the `INSTANCE_CRED` variable.

You can find the instance file in the following location:

```
/etc/tivoli/tsm/instanceList.obj
```

5. Reinstall Tivoli Storage Manager. See Chapter 2, “Installing the server components,” on page 49.

If the `instanceList.obj` file does not exist, you need to recreate your server instances using the following steps:

- a. Recreate your server instances. See “Creating the server instance” on page 62.

Tip: The installation wizard configures the server instances but you must verify that they exist. If they do not exist, you must manually configure them.

- b. Catalog the database. Log in to each server instance as the instance user, one at a time, and issue the following commands:

```
db2 catalog database tsmdb1
db2 attach to instance_name
db2 update dbm cfg using dftdbpath instance_directory
db2 detach
```

- c. Verify that the server instance was created successfully. Issue this command:

Uninstalling Tivoli Storage Manager

```
/opt/tivoli/tsm/db2/instance/db2ilist
```

- d. Verify that Tivoli Storage Manager recognizes the server instance by listing your directories. Your home directory appears if you did not change it. Your instance directory does appear if you used the configuration wizard. Issue this command:

```
db2 list database directory
```

If you see TSMDB1 listed, you can start the server.

Uninstalling IBM Installation Manager

You can uninstall IBM Installation Manager if you no longer have any products that were installed by IBM Installation Manager.

Before you begin

Before you uninstall IBM Installation Manager, you must ensure that all packages that were installed by IBM Installation Manager are uninstalled. Close IBM Installation Manager before you start the uninstall process.

To view installed packages, issue the following command from a command line:

```
cd /opt/IBM/InstallationManager/eclipse/tools  
./imcl listInstalledPackages
```

Procedure

To uninstall IBM Installation Manager, complete the following steps:

1. Open a command line and change directories to `/var/ibm/InstallationManager/uninstall`.
2. Issue the following command:

```
./uninstall
```

Restriction: You must be logged in to the system as the root user ID.

Part 2. Appendixes

Appendix A. Installation log files

If you experience errors during installation, these errors are recorded in log files that are stored in the IBM Installation Manager logs directory.

You can view installation log files by clicking **File > View Log** from the Installation Manager tool. To collect these log files, click **Help > Export Data for Problem Analysis** from the Installation Manager tool.

Appendix B. System resource requirements

Semaphores, shared memory, and processes are three HP-UX system resources IBM Tivoli Storage Manager uses that might require special configuration and tuning.

Estimating required semaphore resources

IBM Tivoli Storage Manager uses semaphore resources to control its internal operations.

About this task

To estimate the total number of semaphores that the server may need, use the following formula:

$$\text{semaphores} = 60 + (2 \times \text{maxSessions})$$

Where *maxSessions* is the maximum number of concurrent client sessions.

For example, if you expect to have up to 15 client sessions active at the same time, Tivoli Storage Manager needs approximately 90 semaphores.

Note: If you have other applications that use semaphores, you must account for their requirements also when adjusting your kernel configuration.

After you have estimated the number of semaphores, ensure that your kernel configuration contains the correct value. See “Viewing and modifying the kernel configuration” on page 108 for details.

Estimating required process resources

IBM Tivoli Storage Manager uses standard HP-UX processes for concurrent server operations.

About this task

To estimate the total number of processes that the server may need, you can use the following formula:

$$\text{processes} = 60 + (2 \times \text{maxSessions})$$

Where *maxSessions* is the maximum number of concurrent client sessions.

For example, assume that you will have up to 15 client sessions active at the same time. You can calculate that IBM Tivoli Storage Manager needs approximately 90 processes to control its internal operations. You will also need to account for all of the other HP-UX processes that might be running concurrently on your system when computing the total requirements for your HP-UX kernel resources.

After you have estimated the required number of processes, ensure that your kernel configuration contains the correct value. See “Viewing and modifying the kernel configuration” on page 108 for details.

Estimating required number of threads per process

The HP-UX default setting for the maximum number of threads allowed in each process is 64 threads.

When Tivoli Storage Manager is running a high workload or participating in LAN-free data movement, this setting might be too low. To prevent thread creation errors in the Tivoli Storage Manager server, increase the HP-UX maximum number of threads per process to 500.

See “Viewing and modifying the kernel configuration” for details.

Viewing and modifying the kernel configuration

To view or modify your existing kernel configuration, use either the SAM utility program on HP-UX Version 11iv2 or System Management Homepage (SMH) on HP-UX Version 11iv3, or edit the configuration file directly. Base the kernel values on the recommendations of the DB2OSCONF utility.

About this task

Start either SAM or SMH, then select:

1. Kernel Configuration
2. Configurable Parameters

A list of parameters, whose values you can change, is displayed. The list includes:

- **semmns** The maximum number of semaphores
- **shmmax** The maximum amount of available shared memory
- **nproc** The maximum number of processes
- **max_thread_proc** The maximum number of threads allowed in each process

See your HP-UX documentation for details about changing configurable kernel parameters.

Appendix C. Accessibility features for the Tivoli Storage Manager product family

Accessibility features help users who have a disability, such as restricted mobility or limited vision to use information technology products successfully.

Accessibility features

The IBM Tivoli Storage Manager family of products includes the following accessibility features:

- Keyboard-only operation using standard operating-system conventions
- Interfaces that support assistive technology such as screen readers

The command-line interfaces of all products in the product family are accessible.

Tivoli Storage Manager Operations Center provides the following additional accessibility features when you use it with a Mozilla Firefox browser on a Microsoft Windows system:

- Screen magnifiers and content zooming
- High contrast mode

The Operations Center and the Tivoli Storage Manager server can be installed in console mode, which is accessible.

The Operations Center help system is enabled for accessibility. For more information, click the question mark icon on the help system menu bar.

Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

IBM and accessibility

See the IBM Human Ability and Accessibility Center (<http://www.ibm.com/able>) for information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
United States of America*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who want to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

Portions of this code are derived from IBM® Corp. Sample Programs.

© Copyright IBM® Corp. _enter the year or years_. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights Except as expressly granted in this permission, no other permissions,

licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

A glossary is available with terms and definitions for the IBM Tivoli Storage Manager family of products.

See Tivoli Storage Manager glossary.

To view glossaries for other IBM products, see IBM Terminology.

Index

A

- access rights
 - setting
 - before server startup 71
- accessibility features 109
- activating
 - server 71
- active log
 - space requirements 32
 - storage technology selection 20
- administrative commands
 - HALT 75
 - REGISTER LICENSE 75
- API 68
- API configuration 68
- archive failover log space
 - description 43
- archive log
 - space requirements 32
 - storage technology selection 20
- archive log directory 59
- automatic starting, server 73

B

- BACKUP DB command 68
- backups
 - database 75

C

- capacity planning
 - database space requirements
 - estimates based on number of files 28
 - estimates based storage pool capacity 30
 - starting size 28
 - recovery log space requirements
 - active and archive logs 32
 - active log mirror 43
- client nodes
 - reverting to previous server version
 - data affected 91
- client options
 - for shared memory communications 65
- commands
 - administrative, SET DBRECOVERY 75
 - DSMSERV FORMAT 67
- commands, administrative
 - HALT 75
 - REGISTER LICENSE 75
- communication methods
 - Shared Memory 65
 - SNMP 66
 - TCP/IP 64
- compatibility, server with other DB2 products 25
- components
 - installable v
- configuring 57, 61, 62
- configuring, manually 61, 62
- configuring, server instance 61

- configuring, wizard 61
- Console language support 53
- console mode 51
- create server instance 57, 61

D

- data deduplication
 - effects when reverting to previous server version 91
- database
 - backups 75
 - installing 67
 - name 46
 - storage technology selection 20
- database directories 59
- database manager 31, 68
- DB2 commands 95
- DB2 directories 48
- DB2 products, compatibility with the server 25
- db2icrt command 62
- db2profile 73
- default installation directories 48
- device driver, Tivoli Storage Manager v
- directories
 - DB2 48
 - default installation 48
 - devices 48
 - languages 48
 - naming for server 46
- directories, instance 59
- disability 109
- DISK device class
 - checklist for disk systems 18
 - storage technology selection 20
- disk performance
 - checklist for active log 11
 - checklist for server database 9
 - checklist for server recovery log 11
 - checklist for storage pools on disk 18
- disk space 22
- disk systems
 - checklist for active log 11
 - checklist for server database 9
 - checklist for server recovery log 11
 - classification 20
 - selecting 20
 - storage pools on disk 18
- DSMSERV FORMAT command 67
- dsmserv.v6lock 75

E

- enabling communications 64
- estimating
 - threads per process (HP-UX) 108
- expiration
 - server option 71

F

- FILE device class
 - checklist for disk systems 18
 - storage technology selection 20
- files
 - dsmerv.opt.smp 64
- first steps 57
- fix packs 79
- fixes 49

G

- group 59

H

- HALT command 75
- halting the server 75
- hardware requirements
 - Tivoli Storage Manager 22
- home directory 62
- HP-UX
 - system requirements 22

I

- IBM Installation Manager 26
 - uninstalling 102
- IBM Knowledge Center vi
- IBM Tivoli Storage Manager
 - uninstalling 99
 - in silent mode 100
 - using a graphical installation wizard 99
 - using command line in console mode 100
- installable components v
- installation log 50, 51
- Installation Manager 26
 - logs directory 105
- installation packages 49
- installation wizard 50
- installing
 - database 67
 - device support 49
 - fix packs 79
 - graphical user interface
 - using 50
 - minimum requirements for 22
 - recovery log 67
 - server 3, 49
 - using command line in console mode
 - using 51
 - what to know before 3
- installing the server
 - silently 52
- installing the Tivoli Storage Manager server 52
- instance directories 59
- instance user ID 46
- interim fix 79

K

- kernel configuration, viewing and modifying 108
- kernel parameters, tuning
 - db2osconf 58
 - overview 57

- kernel parameters, tuning (*continued*)
 - suggested minimum values 59
 - updating 58
- keyboard 109
- KILL command 75
- Knowledge Center vi

L

- LANGUAGE option 53, 54
- language package 55
- language packages 53
- language support 54
- languages
 - set 54
- license, Tivoli Storage Manager 75
- licenses
 - installable package v
- log files
 - installation 105

M

- maintenance updates 79
- memory requirements 22
- monitoring
 - logs 77
- multiple DB2 copies 25
- multiple servers
 - upgrading
 - multiple servers 76

N

- names, best practices
 - database name 46
 - directories for server 46
 - instance user ID 46
 - server instance 46
 - server name 46
- new features vii

O

- offering 26
- Operations Center v
- options
 - starting the server 71
- options file
 - editing 64
- options, client
 - SSLTCPADMINPORT 65
 - SSLTCPPOINT 65
 - TCPADMINPORT 65
 - TCPPOINT 65
 - TCPWINDOWSIZE 65

P

- package 26
- package group 26
- Passport Advantage 49
- performance
 - configuration best practices 21
 - user limits, setting for optimal performance 71

- planning, capacity
 - database space requirements
 - estimates based on number of files 28
 - estimates based storage pool capacity 30
 - starting size 28
 - recovery log space requirements
 - active log mirror 43
 - recovery log space requirementsv 32
- process resources, server
 - estimating total number 107
 - using SAM with 11iv2 to view or change 108
 - using SMH with 11iv3 to view or change 108
- publications vi

R

- recovery log
 - archive failover log space 43
 - installing 67
- reference, DB2 commands 95
- REGISTER LICENSE command 75
- repository 26
- resource for HP system
 - process considerations 107
 - semaphore considerations 107
- reverting to previous server version 91
- rollback 45

S

- scripts
 - rc.dsmserv 73
 - starting servers automatically 73
- Secure Sockets Layer (SSL) 64
 - communication using 66
 - Transport Layer Security (TLS) 66
- semaphore resources, server
 - estimating number of 107
 - using SAM with 11iv2 to view or change 108
 - using SMH with 11iv3 to view or change 108
- server
 - after upgrade
 - reverting to previous server version 91
 - before upgrade
 - importance of preparation steps 91
 - compatibility
 - DB2 products 25
 - naming best practices 46
 - performance optimization 3
 - starting
 - automatic 73
 - stopping 75
 - upgrading
 - V6 to 7.1.7 83
 - V6.1 to V7.1.7 90
 - V6.2 to V7.1.7 84
 - V6.3 to V7.1.7 84
- server active log
 - checklist for disks 11
- server archive log
 - checklist for disks 11
- server database
 - checklist for disks 9
 - directories 9
 - reorganization options 70
 - storage paths 9

- server hardware
 - checklist for server system 4
 - checklist for storage pools on disk 18
 - storage technology choices 20
- server instance 61, 62
- server instance, creating 62
- server instances
 - naming 46
 - naming best practices 46
- server license 75
- server options
 - dsmserv.opt.smp 64
 - tailoring 64
- server options file
 - setting 64
- server recovery log
 - checklist for disks 11
- server,
 - activating 71
 - setting up 71
 - starting 71
- server, Tivoli Storage Manager
 - halting 75
 - options 64
- SET DBRECOVERY 75
- shared memory client options 65
- shared memory communications method 65
- shared resources directory 26
- silent installation
 - Tivoli Storage Manager 52
- SNMP communication method 66
- software requirements
 - Tivoli Storage Manager 22
- SSL (Secure Sockets Layer)
 - communication using 66
 - Transport Layer Security 66
- SSLTCPADMINPORT option 65
- SSLTCPPOPT option 65
- starting
 - server 71
- starting server
 - from user ID 73
- starting servers automatically 73
- stopping
 - server 75
- storage pools 18
 - reverting to previous server version 91
 - storage technology selection 20
- storage technology selection 20
- summary of amendments
 - Version 7.1.7 vii
- system resource requirements for server 107

T

- TCP/IP
 - setting options 64
 - Version 4 64
 - Version 6 64
- TCPNODELAY option 65
- TCPPOPT option 65
- TCPWINDOWSIZE option 65
- technical changes vii
- temporary disk space 31
- temporary space 31
- threads 108

- time
 - server upgrade 84
- Tivoli Monitoring for Tivoli Storage Manager
 - uninstalling IBM Installation Manager 102
- Tivoli Storage Manager
 - installation 50, 51
 - installation packages 49
 - server changes
 - Version 7.1.7 vii
 - uninstalling 99
 - in silent mode 100
 - using a graphical installation wizard 99
 - using command line in console mode 100
 - upgrading
 - V6 to 7.1.7 83
 - V6.1 to V7.1.7 90
 - V6.2 to V7.1.7 84
 - V6.3 to V7.1.7 84
- Tivoli Storage Manager device driver, installable package v
- Tivoli Storage Manager fix packs 79
- Tivoli Storage Manager support site 49
- Tivoli Storage Manager, setting up 71
- translation features 53
- translations 53
- Transport Layer Security (TLS) 66

U

- ulimits
 - setting
 - before server startup 71
- Uninstall
 - IBM Installation Manager 102
- uninstalling 101
- uninstalling and reinstalling 101
- updating 55
- upgrade
 - server
 - estimated time 84
 - V6 to 7.1.7 83
 - V6.1 to V7.1.7 90
 - V6.2 to V7.1.7 84
 - V6.3 to V7.1.7 84
- US English 54
- user ID 59
- user limits 71
 - setting
 - before server startup 71

W

- wizard 57
- worksheet
 - server space planning 27



Product Number: 5608-E01
5608-E02
5608-E03

Printed in USA