

Symantec Backup Exec 2010 Tuning and Performance Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Documentation version: 2010

Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

Symantec and the Symantec Logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Contents

Chapter 1	Analyzing the backup environment	5
	How to analyze your backup requirements	5
	Factors to consider about disk or tape	9
	About designing a backup environment	11
	Calculate the required data transfer rate for backups	12
	Calculate how long it takes to back up to tape	14
	Why faster tape drives are not always better	15
	Calculate how many tape drives are needed	15
	Calculate how long it takes to back up to disk	17
	Calculate the required data transfer rate over a network	17
	Calculate the size of the Backup Exec catalog	19
	About sizing the Backup Exec catalog	22
	Calculate the number of tapes needed for full and incremental backups	22
	Calculate the size of the tape library needed to store backups	25
Chapter 2	Measuring performance	27
	About measuring Backup Exec performance	27
	Controlling system variables for consistent testing conditions	28
	Running a performance test	28
	Controlling network variables during performance testing	28
	Controlling remote agent variables during performance testing	29
	Controlling data variables during performance testing	29
	Evaluating performance through the Job Monitor	30
	Evaluating Windows system components	31
	Monitoring the Windows CPU load	32
	Monitoring memory use	32
	Monitoring disk load	33
Chapter 3	Tuning the media server	35
	Tips for backing up many small files	35
	Use backup-to-disk folders	36
	Use the Deduplication Option	37
	Tips for including and excluding selections	37

	Limit the backup frequency	37
	Run verify operations outside of the backup window	37
	Use disk staging	38
	Keep media servers upgraded	38
	Increase disk performance	38
Chapter 4	Tuning the Backup Exec data transfer path	41
	About tuning the data transfer path	41
	Tuning suggestions for the data path	42
	Remote agent performance	43
	Network performance	44
	Media server performance	46
	About the Windows data buffer size	49
	Storage device performance	49
Appendix A	Additional Resources	53
	Additional Resources	53

Analyzing the backup environment

This chapter includes the following topics:

- How to analyze your backup requirements
- Factors to consider about disk or tape
- About designing a backup environment
- Calculate the required data transfer rate for backups
- Calculate how long it takes to back up to tape
- Why faster tape drives are not always better
- Calculate how many tape drives are needed
- Calculate how long it takes to back up to disk
- Calculate the required data transfer rate over a network
- Calculate the size of the Backup Exec catalog
- About sizing the Backup Exec catalog
- Calculate the number of tapes needed for full and incremental backups
- Calculate the size of the tape library needed to store backups

How to analyze your backup requirements

To estimate your backup requirements, you must consider your environment. Many performance issues are linked to hardware or environmental issues. An

understanding of the entire backup data path is important to determine the maximum performance that you can expect from your installation.

Every backup environment has a bottleneck. It may be a fast bottleneck, but it still determines the maximum performance that you can obtain.

How can you configure Backup Exec or the environment to increase backup performance? Many elements influence your backup strategy. Analyze these elements and then make backup decisions according to your site's priorities.

Table 1-1 How to analyze backup requirements

Decisions	Considerations
Which computers do you have to back up?	Identify all computers that you want to back up, and list the following information for each computer: <ul style="list-style-type: none">■ Name■ Operating system■ Database types and versions■ Network technology, such as 1000BaseT■ Any attached disk drives, tape drives, or robotic libraries■ Model type of each drive or library■ Any applications that you want to back up, such as Oracle, DB2, or Exchange

Table 1-1 How to analyze backup requirements (*continued*)

Decisions	Considerations
<p>How much data do you have to back up?</p>	<p>Calculate how much data you need to back up. Include the total disk space on each individual computer, including the space that the databases use. Remember to add the space on mirrored disks only once.</p> <p>If you plan to back up databases, identify the database engines, their version numbers, and the method to back them up. Backup Exec can back up several database engines and raw file systems, and databases can be backed up while they are online or offline. To back up any database while it is online, you need a Backup Exec database agent.</p> <p>With a snapshot backup of databases using raw partitions, you back up as much data as the total size of the raw partition. Remember to add the size of the database backups to the final calculations when calculating how much data must be backed up.</p> <p>If you plan to back up application servers such as Exchange or Lotus Notes you must identify their types and application release numbers. You may need a Backup Exec agent to properly back up an application server.</p>
<p>What types of backups do you need and how often do you need them?</p>	<p>To properly size your backup environment, you must decide on the type and frequency of your backups. Will you perform daily incremental and weekly full backups? Monthly or bi-weekly full backups?</p> <p>The frequency of the backups has a direct effect on the following items:</p> <ul style="list-style-type: none"> ■ Disk or tape requirements ■ Data transfer rate considerations ■ Restore opportunities <p>See “Factors to consider about disk or tape” on page 9.</p>

Table 1-1 How to analyze backup requirements (*continued*)

Decisions	Considerations
How much time is available to run each backup?	<p>The length of the backup window dictates several aspects of your backup strategy. For example, you may want a larger window to back up multiple, high-capacity servers. Or you may consider the use of advanced Backup Exec features such as synthetic backups, or a local snapshot method.</p>
How long should you retain backups?	<p>The amount of time a backup is kept is known as the retention period. A common retention period is one month for incremental backups and six months for full backups. With these retention periods, you can restore any daily file change from the previous month and restore data from full backups for the previous six months.</p> <p>The length of the retention period depends on your business needs. However, the length of the retention period is directly proportional to the number of tapes you need and the size of the Backup Exec catalog database. The Backup Exec catalog database keeps track of all the information on all your disk drives and tapes. The catalog size is relative to the retention periods and to the frequency of your backups.</p> <p>Also, database management services may become bottlenecks.</p>
If backups are sent off site, how long must they remain off site?	<p>If you plan to send tapes off site, identify which tapes to send off site and how long they remain off site. You might decide to duplicate all full backups, or only a select few. You might also decide to duplicate certain computers and exclude others. As tapes are sent off site, you must buy new tapes to replace them until they are recycled back from off site storage. If you forget this detail, you may run out of tapes when you most need them.</p>

Table 1-1 How to analyze backup requirements (*continued*)

Decisions	Considerations
What is your network technology?	If you plan to back up any computers over a network, note the network types. Based on the amount of data that you want to back up and the frequency of those backups, consider installing a private network for backups.
What new computers do you expect to add in the next six months?	Plan for future growth when you design your backup environment. Analyze the potential growth of your environment to ensure that your current backup solution can accommodate future requirements. Add any resulting growth factor that you incur to your total backup solution.
What are the types of data that you want to back up?	Consider the data types that you want to back up, such as text, graphics, or databases. How compressible is the data? How many files are involved? Will the data be encrypted? Note that encrypted backups may run slower.
Is the data local or remote?	Consider the characteristics of the storage subsystem. What is the exact data path? How busy is the storage subsystem?
Are you prepared to test component changes?	Because hardware and software infrastructure change over time, create an independent test-backup environment to ensure your production environment works with the changed components.

See [“About designing a backup environment”](#) on page 11.

Factors to consider about disk or tape

Disk is now a common backup medium. Backup data that is on disk generally provides faster restores.

You can tune disk-based storage for performance the same way that you can tune tape-based storage. The optimal buffer settings for a site can vary according to its configuration. It takes thorough testing to determine these settings.

Disk-based storage can be useful if you have a lot of incremental backups and the percentage of data change is small. If the volume of data is insufficient in incremental copies to ensure efficient writing to tape, consider disk storage. After writing the data to disk, you can use a duplicate backup data job to copy the data to tape. This arrangement can produce faster backups and prevent wear and tear on your tape drives.

Consider the following factors when backing up data to disk or tape:

- Will you use short or long retention periods?
Disk is well suited for short retention periods; tape is better suited for longer retention periods.
- Will you use intermediate staging or long-term storage?
Disk is suited for staging; tape for long-term storage.
- Do you plan to run incremental or full backups?
Disk is better suited to low volume incremental backups.
- Do you plan to run synthetic backups?
Synthetic backups eliminate the need to perform recurring full backups for supported remote resources.
- Is data recovery time important?
Restore from disk is usually faster than from tape.
- What is the speed of the backups?
If backups are too slow to keep the tape in motion, send the backups to disk. Later, a duplicate backup data job (staging) can move the backup images to tape.
- What are the sizes of the backups?
If the backups are too small to keep the tape in motion, send the backups to disk. Small backups may include incrementals and frequent backups of small database log files. A duplicate backup data job can later move the backup images to tape.

The following are some benefits of backing up to disk rather than tape:

- Faster access to data
Most tape drives have a "time to data" of close to two minutes. Time is required to move the tape from its slot, load it into the drive, and seek to an appropriate place on the tape. Disk has an effective time to data of 0 seconds. Restoring a large file system from 15 different tapes can add almost one hour to the restore: a two-minute delay per tape for load and seek, and a possible two-minute delay per tape for rewind and unload.
- Fewer full backups

With tape-based systems, full backups must be done regularly because of the "time to data" issue. If full backups are not done regularly, a restore may require too many tapes from incremental backups. As a result, the time to restore increases, as does the chance that a single tape may cause the restore to fail.

- Use of the Granular Recovery Technology (GRT) feature without having to stage data from tape to disk

You can use GRT to restore certain individual items from backup sets. For example, you can restore an email message from a backup without having to restore the entire mailbox. If you send the GRT-enabled backup to a backup-to-disk folder rather than to tape, you can restore individual items without staging the data from tape to disk.

- Use of the Deduplication Option

You can back up to a disk-based deduplication storage folder on the media server. Deduplication reduces storage requirements and network traffic.

About designing a backup environment

After an analysis of your backup requirements, you can begin designing your backup environment.

Table 1-2 About designing a backup environment

Step	Description	More information
1.	Calculate the rate of transfer that Backup Exec must achieve to complete a backup of all your data in the allowed time window.	See "Calculate the required data transfer rate for backups" on page 12.
2.	Decide what kind of tape drive technology meets your needs by calculating how long it takes to back up to tape.	See "Calculate how long it takes to back up to tape" on page 14.
3.	Calculate how many tape drives are needed.	See "Calculate how many tape drives are needed" on page 15.
4.	Calculate how long it takes to back up to disk.	See "Calculate how long it takes to back up to disk" on page 17.

Table 1-2 About designing a backup environment (*continued*)

Step	Description	More information
5	Calculate how long it takes to move data over the network, if necessary. You must move data from the remote computer to the media server fast enough to finish backups within your backup window.	See “Calculate the required data transfer rate over a network” on page 17.
6	Calculate how much disk space is needed to store the Backup Exec catalog.	See “About sizing the Backup Exec catalog” on page 22.
7.	Calculate how many tapes are needed to store and retrieve your backups.	See “Calculate the number of tapes needed for full and incremental backups” on page 22.
8.	Calculate how many robotic library tape slots are needed to store all your backups.	See “Calculate the size of the tape library needed to store backups” on page 25.

Calculate the required data transfer rate for backups

You can calculate the rate of transfer that Backup Exec must achieve to complete a backup of all your data in the allowed time window. Use the following formula to calculate the minimum average data transfer rate for full and incremental backups:

$$\text{Minimum average data transfer rate} = (\text{Amount of data to back up}) / (\text{Backup window})$$

On average, the daily change in data for many systems is between 10 and 20 percent. Calculate a change of 20% in the amount of data to back up. Then divide it by the backup window for the backup rate for incremental backups.

If you run incremental backups, account for data that undergoes changes. For example, if the same 20% of the data changes daily, your incremental backup is smaller than if a different 20% changes every day.

For example, assume that you back up 500 GB during a full backup and that the daily backup window is eight hours. The minimum average data transfer rate for 500 GB per eight hours is 62.5 GB per hour. Assume that you back up 100 GB

during an incremental backup with the daily backup window. The minimum average transfer rate for 100 GB per eight hours is 12.5 GB per hour.

For your minimum average weekend transfer rate, divide the amount of data that you must back up by the length of the weekend backup window.

For example, assume that the amount of data to back up during a full backup is 1000 gigabytes. Assume also that the daily backup window is eight hours.

The formula for the minimum average transfer rate is as follows:

$$(\text{Data}/(\text{Backup window})) = 1000 \text{ gigabytes per eight hours} = 125 \text{ gigabytes per hour}$$

Table 1-3 Examples for calculating the required data transfer rate for backups

Example	Factors	Solution
Example 1	One LTO-4 tape drive with a drive transfer rate of 432 GB per hour:	Actual data transfer rate = 1000 GB/((one drive) * (432 GB per hour)) = 2.31 hours With a data transfer rate of 432 GB per hour, a single LTO-4 tape drive takes 2.31 hours to perform a 1000-GB backup.
Example 2	One LTO-5 tape drive with a drive transfer rate of 504 GB per hour:	Actual data transfer rate = 1000 gigabytes/((one drive) * (504 GB per hour)) = 1.98 hours With a data transfer rate of 504 GB per hour, a single LTO-5 tape drive takes 1.98 hours to perform a 1000-GB backup.

Depending on the several factors that can influence the transfer rates of your tape drives, you can obtain higher or lower transfer rates. These example solutions are approximations of what you can expect. A backup of encrypted data may take more time.

Calculate how long it takes to back up to tape

When you know your ideal data transfer rates for backups, you can calculate what kind of tape drive technology meets your needs. With the length of the backup windows and the amount of data to back up, you can calculate the number of required tape drives.

Table 1-4 Tape drive data transfer rates

Drive	GB per hour (no compression)	GB per hour (2:1 compression)	Native MB per second	Native MB per minute
LTO-1	54	108	15	900
LTO-2	105	216	40	2400
LTO-3	288	576	80	4800
LTO-4	432	864	120	7200
LTO-5	504	1008	140	8400
SDLT 320	57	115	16	960
SDLT 600	129	259	36	2160
STK 9940B	108	252 (2.33:1)	30	1800
STK T10000	422	844	120	7200
DAT 320	43	86	12	720
DAT 160	24	49	6.9	414

The values are those published by their individual manufacturers and observed in real-life situations. Keep in mind that device manufacturers list optimum rates for their devices.

In reality, it is rare to achieve those values when a computer has to deal with the following issues:

- The overhead of the operating system
- CPU loads
- Bus architecture
- Data types
- Other hardware and software issues

- File system directory overhead
- Average file sizes (big files stream faster)

The GB per hour values in the table above represent transfer rates for several devices, with and without compression. When you design your backup system, consider the nature of both your data and your environment. Estimate on the conservative side when planning capacity.

To calculate the length of your backups using a particular tape drive, use the following formula:

Actual data transfer rate = (Amount of data to back up)/((Number of drives) * (Tape drive transfer rate))

See [“Calculate how many tape drives are needed”](#) on page 15.

Why faster tape drives are not always better

The fastest tape technology may not be the most appropriate for your site. Tape drives also have a minimum speed below which they start to operate inefficiently. This figure is known as the minimum streaming speed and is usually around 40% of the native (no compression) speed of the device. If the drive receives data at less than minimum streaming speed, it operates in a stop-and-start mode (shoe shining). In this mode the drive empties the data buffers faster than they can be filled and has to stop while the buffers fill up again. When the buffers are full the drive must start and then reposition the tape before it writes the next data block. This stop-and-start behavior can cause excessive wear to both the tape and the drive and also results in a further slowing down of the backup. For this reason, the fastest tape technology may not always be the most appropriate one to use.

Unlike tape drives, disk devices (including VTLs) do not have a minimum streaming speed. A good strategy may be to stage slower backups to disk before you duplicate them to tape. The duplication of the backup image from disk runs faster than the original slow backup.

See [“Use disk staging”](#) on page 38.

Calculate how many tape drives are needed

The formula to calculate how many tape drives are needed is as follows:

Number of drives = (Amount of data to back up)/((Backup window) * (Tape drive transfer rate))

For example, assume that the amount of data to back up is 1000 gigabytes, and that the backup window is eight hours.

Table 1-5 Examples for calculating how many tape drives are needed

Example	Factors	Solution
Example 1	This example uses an LTO-2 tape drive with a drive transfer rate of 105 gigabytes per hour.	Number of drives = 1000 gigabytes/ ((Eight hours) * (105 gigabytes per hour)) = .1.19 = 2 drives
Example 2	This example uses an LTO-4 tape drive with a drive transfer rate of 432 gigabytes per hour.	Number of drives = 1000 gigabytes/((Eight hours) * (432 gigabytes per hour)) = 0.29 = 1 drive

You can calculate the number of drives that are needed to perform a backup. The difficulty is the ability to spread the data streams evenly across all drives. To spread the data streams, you can experiment with various backup schedules and your hardware configuration.

To calculate how many tape devices you need, calculate how many tape devices you can attach to a drive controller.

To calculate the maximum number of drives that you can attach to a controller, you need the manufacturers' maximum transfer rates for drives and controllers. Failure to use maximum transfer rates for your calculations can result in saturated controllers and unpredictable results.

In practice, your transfer rates might be slower because of the inherent overhead of several variables. Variables include the file system layout, system CPU load, and memory usage

Table 1-6 Drive controller data transfer rates

Drive Controller	Theoretical MB per second	Theoretical MB per minute	Theoretical GB per hour
Wide Ultra 2 SCSI	80	4800	288
iSCSI (on a 1-gigabit network)	100	6000	360
1 gigabit Fibre Channel	100	6000	360
Ultra-3 SCSI	160	9600	576
SATA/150	187	11250	675

Table 1-6 Drive controller data transfer rates (*continued*)

Drive Controller	Theoretical MB per second	Theoretical MB per minute	Theoretical GB per hour
2 gigabit Fibre Channel	200	12000	720
Ultra320 SCSI	320	19200	1152
SATA/300	375	22500	1350
3 gigabit SAS	375	22500	1350
4 gigabit Fibre Channel	400	24000	1440
6 gigabit SAS	750	45000	2700
8 gigabit Fibre Channel	850	51000	3060
iSCSI (on a 10-gigabit network)	1250	75000	4500

Calculate how long it takes to back up to disk

Disk performance is more difficult to predict than data transfer rates for tapes. Backup speed depends on the type of disk and on the disk layout. In general, the speed of the backup depends on the speed of the disks and disk controllers that the backup is written to.

Calculate the required data transfer rate over a network

To back up over a network, you must move data from the remote computer to the media server fast enough to finish backups within your backup window. Calculate the data transfer rates for your networks to help you identify potential bottlenecks. Several solutions are available for dealing with multiple networks and bottlenecks.

To calculate the required data transfer rate, use the following formula:

Required network data transfer rate = (Amount of data to back up) / (Backup window)

The following table shows the typical transfer rates of some common network technologies.

Table 1-7 Network data transfer rates

Network Technology	Theoretical GB per hour	Typical GB per hour
100BaseT (switched)	36	25
1000BaseT (switched)	360	250
10000BaseT (switched)	3600	2500

For example, to calculate network transfer rates, assume that the amount of data to back up is 500 gigabytes with a backup window of eight hours.

The formula is as follows:

Required network transfer rate = 500 GB/8 hours = 62.5 GB per hour

Table 1-8 Examples for calculating the required data transfer rate over a network

Example	Factors	Solution
Example 1	This example uses the 100BaseT (switched) network technology with a typical transfer rate of 25 GB per hour.	<p>A single 100BaseT network has a transfer rate of 25 GB per hour. This network cannot handle your required data transfer rate of 62.5 GB per hour.</p> <p>In this situation, you may want to review other options, such as the following:</p> <ul style="list-style-type: none"> ■ Back up your data over a faster network (1000BaseT). ■ Back up large servers to dedicated tape drives. ■ Back up over SAN connections by using the SAN Shared Storage Option. ■ Perform off-host backups to present a snapshot directly to a media server. ■ Perform your backups during a longer time window. ■ Perform your backups over faster dedicated networks.
Example 2	This example uses the 1000BaseT (switched) network technology with a typical transfer rate = 250 GB per hour.	A single 1000BaseT network has a transfer rate of 250 GB per hour. This network technology can handle your backups with room to spare.

Calculate the size of the Backup Exec catalog

When you design your backup environment, calculate how much disk space is needed to store the Backup Exec catalog. The catalog tracks all of the files that are backed up.

The catalog's size depends on the following variables, for both full backups and incremental backups:

- The number of files being backed up.
- The frequency and the overwrite period of the backups.

You can use either of two methods to calculate the Backup Exec catalog size. In both cases, since data volumes grow over time, you should factor in expected growth when calculating total disk space used.

To use the first method, you must know the following details:

- The number of files that are held online.
- The number of full and incremental backups that are retained at any time.

To calculate the size in GB for a particular backup, use the following formula:

Catalog size = (135 * Number of files in all backups) / 1 GB

The number 135 is based on the average file name length of 30 characters.

Determine the approximate number of copies of each file that is held in backups and the typical file size.

Use the following formula to estimate the number of copies:

Number of copies of each file that is held in backups = Number of full backups + 10% of the number of incremental backups held

The second method for calculating the catalog size is to multiply the total amount of data in the production environment by a small percentage (such as 2%). The total amount of data in the production environment is not the total size of all backups. Also, note that 2% is only an example; the following formulas help you calculate a percentage that is appropriate for your environment.

Use this method only in environments in which it is easy to determine the typical file size, typical retention periods, and typical incremental change rates. In some cases, the catalog size that is obtained using this method may vary significantly from the eventual catalog size.

To use this method, you must determine the approximate number of copies of each file that is held in backups and the typical file size. Use the following formula to estimate the number of copies:

Number of copies of each file that is held in backups = Number of full backups + 10% of the number of incremental backups held

To calculate the multiplying percentage, use the following formulas:

- Multiplying percentage = (120 * Number of files that are held in backups / Average file size) * 100%

- $\text{Multiplying percentage} = \frac{132 * \text{Number of files that are held in backups}}{\text{Average file size}} * 100\%$

Then, you can estimate the size of the catalog by using the following formula:

$\text{Size of the catalog} = \text{Total disk space used} * \text{Multiplying percentage}$

Table 1-9 Examples for calculating the size of the Backup Exec catalog

Example	Factors	Solution
Method 1	<p>Assume the following:</p> <ul style="list-style-type: none"> ■ Number of full backups per month: four ■ Retention period for full backups: six months ■ Total number of full backups retained: 24 ■ Number of incremental backups per month: 25 ■ Total number of the files that are held online (total number of files in a full backup): 17,500,000 	<p>Number of copies of each file retained: $24 + (25 * 10\%) = 26.5$</p> <p>Catalog size for each file retained: $(120 * 26.5 \text{ copies}) = 3180 \text{ bytes}$</p> <p>Total catalog space required: $(3180 * 17,500,000 \text{ files}) / 1\text{GB} = 54 \text{ GB}$</p>
Method 2	<p>Assume the following:</p> <ul style="list-style-type: none"> ■ Number of full backups per month: 4 ■ Retention period for full backups: six months ■ Retention period for full backups: six months ■ Retention period for full backups: six months ■ Average file size: 70 KB ■ Total disk space that is used on all servers in the domain: 1.4 TB 	<p>Number of copies of each file retained: $24 + (25 * 10\%) = 26.5$</p> <p>Backup Exec catalog size for each file retained: $(120 * 26.5 \text{ copies}) = 3180 \text{ bytes}$</p> <p>$(120 * 26.5 \text{ copies}) = 3180 \text{ bytes}$</p> <p>Multiplying percentage: $(3180 / 70000) * 100\% = 4.5\%$</p> <p>$(3498 / 70000) * 100\% = 5\%$</p> <p>Total catalog space required: $(1,400 \text{ GB} * 4.5\%) = 63 \text{ GB}$</p>

About sizing the Backup Exec catalog

The size of the catalog depends on the number of files in the backups and the number of copies of the backups that are retained. As a result, the catalog has the potential to grow quite large. When you estimate the ultimate size of the catalog, consider if it can be backed up in an acceptable time window, and if the general housekeeping activities can complete within their execution windows.

In general, Symantec recommends that you plan your environment to meet the following criteria:

- The amount of data that is held in the online catalog should not exceed 750 GB. Use catalog archiving to keep the online portion of the catalog under 750 GB.
- The total number of catalog files should not exceed 200,000. This number equals the total of all retained copies of all backups from all clients held both online and in the catalog archive.

The speed of the storage and the power of the server influence the actual limits of acceptable catalog performance. Your catalog's performance may vary significantly from the guidelines that are provided in this section.

If you expect that your catalog may exceed these limits, consider deploying multiple Backup Exec domains in your environment.

Calculate the number of tapes needed for full and incremental backups

Calculate how many tapes that you need to store and retrieve your backups.

The number of tapes depends on the following:

- The amount of data that you back up.
- The frequency of your backups.
- The planned retention periods.
- The capacity of the media that is used to store your backups.

If you expect your site's workload to increase over time, you can ease the pain of future upgrades by planning for expansion. Design your initial backup architecture so it can evolve to support more clients and servers. Invest in the faster, higher-capacity components that can serve your needs beyond the present.

A formula for calculating your tape needs is as follows:

Number of tapes = (Amount of data to back up) / (Tape capacity)

To calculate how many tapes are needed based on all your requirements, the previous formula can be expanded to the following:

$$\text{Number of tapes} = ((\text{Amount of data to back up}) * (\text{Frequency of backups}) * (\text{Retention period})) / (\text{Tape capacity})$$

Table 1-10 Tape capacities

Drive	Theoretical GB (no compression)	Theoretical GB (2:1 compression)
LTO-1	100	200
LTO-2	200	400
LTO-3	400	800
LTO-4	800	1600
LTO-5	1600	3200
SDLT 320	160	320
SDLT 600	300	600
STK 9940B	200	400
STK T10000	500	1000
STK T10000B	1000	2000

For the examples in the next table, assume the following:

Size of full backups = 500 gigabytes * 4 (per month) * 6 months = 12 terabytes

Size of incremental backups = (20% of 500 gigabytes) * 30 * 1 month = 3 terabytes

Total data tracked = 12 terabytes + 3 terabytes = 15 terabytes

Table 1-11 Examples for calculating the number of tapes needed for full and incremental backups

Example	Factors	Solution
Example 1	This example uses an LTO-2 tape drive without compression (200-GB capacity)	<p>Tapes that are needed for full backups = $12 \text{ terabytes} / 200 \text{ GB} = 60$</p> <p>Tapes that are needed for incremental backups = $3 \text{ terabytes} / 200 \text{ GB} = 15$</p> <p>Total tapes needed = $60 + 15 = 75$ tapes</p>
Example 2	This example uses an LTO-2 tape drive with 2:1 compression (400-GB capacity)	<p>Tapes that are needed for full backups = $12 \text{ terabytes} / 400 \text{ GB} = 30$</p> <p>Tapes that are needed for incremental backups = $3 \text{ terabytes} / 400 \text{ GB} = 7.5$</p> <p>Total tapes needed = $30 + 7.5 = 37.5$ tapes</p>
Example 3	This example uses an LTO-3 tape drive without compression (400-GB capacity)	<p>Tapes that are needed for full backups = $12 \text{ terabytes} / 400 \text{ GB} = 30$</p> <p>Tapes that are needed for incremental backups = $3 \text{ terabytes} / 400 \text{ gigabytes} = 7.5 \approx 8$</p> <p>Total tapes needed = $30 + 8 = 38$ tapes</p> <p>Tape capacity with 2:1 compression = 800 gigabytes:</p>
Example 4	This example uses an LTO-3 tape drive with 2:1 compression (800-GB capacity)	<p>Tapes that are needed for full backups = $12 \text{ terabytes} / 800 \text{ GB} = 15$</p> <p>Tapes that are needed for incremental backups = $3 \text{ terabytes} / 800 \text{ GB} = 3.75 \approx 4$</p> <p>Total tapes needed = $15 + 4 = 19$ tapes</p>

Calculate the size of the tape library needed to store backups

To calculate how many robotic library tape slots are needed to store all your backups, use the number of backup tapes that you calculated in a previous section. Then, add tapes for catalog backup and for cleaning.

The formula is the following:

Tape slots needed = (Number of tapes that are needed for backups) + (Number of tapes that are needed for catalog backups) + 1 (for a cleaning tape)

For tapes for catalog backup, a typical example is 2.

Additional tapes may be needed for the following situations:

- If you plan to duplicate tapes or to reserve some media for special (non-backup) use, add those tapes to the formula.
- Add the tapes that are needed for future data growth. Make sure that your system has a viable upgrade path as new tape drives become available.

Calculate the size of the tape library needed to store backups

Measuring performance

This chapter includes the following topics:

- [About measuring Backup Exec performance](#)
- [Controlling system variables for consistent testing conditions](#)
- [Running a performance test](#)
- [Controlling network variables during performance testing](#)
- [Controlling remote agent variables during performance testing](#)
- [Controlling data variables during performance testing](#)
- [Evaluating performance through the Job Monitor](#)
- [Evaluating Windows system components](#)
- [Monitoring the Windows CPU load](#)
- [Monitoring memory use](#)
- [Monitoring disk load](#)

About measuring Backup Exec performance

You can measure Backup Exec performance by one of the following criteria:

- The length of time in your environment that you can allow for backup operations to complete without affecting the performances of the server, end users, or the network. This length of time is known as the backup window.
- The length of time that is required for a critical restore operation to complete.

You need to use the performance metrics that are more reliable and reproducible than using wall clock time. After establishing accurate metrics, you can measure

the current performance of Backup Exec and your system components to compile a baseline performance benchmark. With a baseline, you can apply changes in a controlled way. By measuring performance after each change, you can accurately measure the effect of each change on Backup Exec performance.

Controlling system variables for consistent testing conditions

For reliable performance evaluation, eliminate as many unpredictable variables as possible to create a consistent backup environment. Only a consistent environment can produce reliable and reproducible performance measurements. This section explains some of the variables to consider as they relate to the media server, the network, the remote agent, or the data itself.

Eliminate all other Backup Exec activity from your environment when you measure the performance of a particular Backup Exec operation.

Running a performance test

Use the following procedure to run a test. This procedure helps prevent the Backup Exec scheduler from running other backups during the test.

To run a performance test

- 1 Before you start the performance test, check the Job Monitor to ensure that there are no Backup Exec jobs in progress.
- 2 Submit the job to run now.
- 3 During the performance test, check for non-Backup Exec activity on the server and try to reduce or eliminate it.
- 4 Check the Backup Exec Job Monitor after the performance test for any unexpected activity that may have occurred during the test, such as a restore job.

Controlling network variables during performance testing

Network performance is key to optimum performance with Backup Exec. Ideally, you should use a separate network for testing to prevent unrelated network activity from skewing the results.

In many cases, a separate network is not available. If not, ensure that non-Backup Exec activity is kept to a minimum during the test. If possible, schedule the test when backups are not active. Even occasional bursts of network activity may be enough to skew the test results. If you share the network with production backups occurring for other systems, you must account for this activity during the test.

Another network variable is host name resolution. Backup Exec depends heavily upon a timely resolution of host names to operate correctly. If you have any delays in host name resolution, try to eliminate that delay by ensuring that the DNS is properly configured. An example of such a delay is a reverse name lookup to identify a server name from an incoming connection from an IP address. You can use the Windows `HOSTS` command or the UNIX command `/etc/hosts` for host name resolution on systems in your test environment.

Controlling remote agent variables during performance testing

Make sure that the remote agent is relatively quiescent during performance testing. A lot of activity, especially disk-intensive activity such as Windows virus scanning, can limit the data transfer rate and skew the test results.

Do not allow another media server to access the remote agent during the test. Backup Exec may attempt to back up the same remote agent to two different servers at the same time. The results of a performance test that is in progress can be severely affected.

Different file systems have different performance characteristics. It may not be valid to compare data throughput on UNIX VxFS or Windows FAT file systems to UNIX NFS or Windows NTFS systems. For such a comparison, factor the difference between the file systems into your performance tests and into any conclusions.

Controlling data variables during performance testing

Monitoring the data that is backed up improves the repeatability of performance testing.

A large set of data generates a more reliable, reproducible test than a small set of data. Startup and shutdown overhead within the Backup Exec operation may skew a performance test that uses a small data set. These variables are difficult to keep consistent between test runs and are therefore likely to produce inconsistent test results. A large set of data minimizes the effect of startup and shutdown times.

Design the data set to represent the makeup of the data in the intended production environment. If the data set in the production environment contains many small

files, the data set for the tests should also contain many small files. A representative data set can more accurately predict the Backup Exec performance that can be expected in a production environment.

The type of data can help reveal bottlenecks in the system. Files that contain non-compressible (random) data cause the tape drive to run at its lower rated speed. As long as the other components of the data transfer path can keep up, you can identify the tape drive as the bottleneck. On the other hand, files containing highly-compressible data can be processed at higher rates by the tape drive when hardware compression is enabled. The result may be a higher overall throughput and may expose the network as the bottleneck.

Some values in Backup Exec provide data amounts in kilobytes and rates in kilobytes per second. For greater accuracy, divide by 1024 rather than rounding off to 1000 when you convert from kilobytes to megabytes or kilobytes per second to megabytes per second.

For best results with controlling data variables, consider the following tips:

- If possible, move the data you use for testing to its own drive or logical partition (not a mirrored drive).
- Defragment the drive before you begin performance testing.
- To test restores, start with an empty disk drive or a recently defragmented disk drive with ample empty space.
- When you test backups to tape, always start each test with an empty piece of media.
- When you test restores from tape, always restore from the same backup image on the tape to achieve consistent results between tests.

Evaluating performance through the Job Monitor

You can obtain statistics for backup or restore operations from the Job Monitor. Wait until all the individual backup operations are complete. The statistics that display in the Job Monitor are for each of the individual backup operations. The statistics do not reflect the total data throughput to the tape drive.

To evaluate performance through the Job Monitor

- 1 Run the backup or restore job.
- 2 On the navigation tab, click **Job Monitor**.
- 3 Verify that the backup or restore completed successfully.
- 4 On the **Job list** tab, in the **Job History** pane, select the completed job that you want to view.

- 5 In the task pane, under **General Tasks**, click **Properties**.
- 6 View the information on the **Job History** tab.
- 7 Obtain the Backup Exec performance statistics from the following fields in the job history:

Job started/Job ended	Displays the time window during which the backup or restore job took place.
Elapsed time	Displays the total elapsed time from when the job was initiated to job completion. It can be used as an indication of total wall clock time for the operation.
Job rate	Displays the amount of data that was backed up per minute for the entire job.
Byte count	Displays the number of bytes that were processed. Compare this value to the job rate.

Evaluating Windows system components

In addition to your evaluation of Backup Exec's performance, you should also verify that common system resources are in adequate supply. You may want to use the Windows Performance Monitor utility. For information about using the Performance Monitor, refer to your Microsoft documentation.

The Performance Monitor organizes information by object, counter, and instance.

An object is a system resource category, such as a processor or physical disk. Properties of an object are counters. Counters for the **Processor** object include **%Processor Time**, which is the default counter, and **Interrupts/sec**. Duplicate counters are handled by instances. For example, to monitor the **%Processor Time** of a specific CPU on a multiple CPU system, the **Processor** object is selected. Then the **%Processor Time** counter for that object is selected, followed by the specific CPU instance for the counter.

In the Performance Monitor, you can view data in real-time format or collect the data in a log for future analysis. Specific components to evaluate include CPU load, memory use, and disk load.

You should use a remote host for monitoring of the test host, to reduce the load that might otherwise skew results.

Monitoring the Windows CPU load

To determine if the system has enough power to accomplish the requested tasks, monitor the **% Processor Time** counter for the **Processor** object to determine how hard the CPU works. Also monitor the **Process Queue Length** counter for the **System** object to determine how many processes actively wait for the processor.

For **% Processor Time**, values of 0 to 80 percent are generally safe. The system is heavily loaded if the values are from 80 percent to 90 percent. Consistent values over 90 percent indicate that the CPU is a bottleneck.

Spikes close to 100 percent are normal and do not necessarily indicate a bottleneck. However, if sustained loads close to 100 percent are observed, consider tuning the system to decrease process load, or upgrade to a faster processor.

Sustained **Processor Queue Lengths** greater than 2 indicate that too many threads wait to be executed. To correctly monitor the **Processor Queue Length** counter, the Performance Monitor must track a thread-related counter. If you consistently see a queue length of 0, verify that a non-zero value can be displayed.

The default scale for the **Processor Queue Length** may not be equal to 1. Be sure to read the data correctly. For example, if the default scale is 10x, then a reading of 40 means that only four processes must wait.

Monitoring memory use

Memory is a critical resource for increasing the performance of backup operations.

When you examine memory usage, view information on the following items:

- **Committed Bytes**

Committed Bytes displays the size of virtual memory that has been committed, as opposed to reserved. Committed memory must have disk storage available or must not require the disk storage because the main memory is large enough. If the number of Committed Bytes approaches or exceeds the amount of physical memory, you may encounter issues with page swapping.

- **Page Faults/sec**

Page Faults/sec is a count of the page faults in the processor. A page fault occurs when a process refers to a virtual memory page that is not in its working set in main memory. A high Page Fault rate may indicate insufficient memory.

Monitoring disk load

To use disk performance counters to monitor the disk performance in Performance Monitor, you may need to enable the counters. Windows may not have enabled the disk performance counters by default for your system.

On Windows 2000, the counters are set by default.

To get more information about disk performance counters

- ◆ Enter the following:

```
diskperf -help
```

To enable the counts and allow disk monitoring

- 1 Enter the following:

```
diskperf -y
```

- 2 Restart the computer.

To disable the counters and cancel disk monitoring

- 1 Enter the following:

```
diskperf -n
```

- 2 Restart the computer.

To monitor disk performance

- 1 Use the **%Disk Time counter** for the **PhysicalDisk** object.

Track the percentage of elapsed time that the selected disk drive services read or write requests.

- 2 Monitor the Avg. Disk Queue Length counter and watch for values greater than 1 that last for more than one second.

Values greater than 1 for more than a second indicate that multiple processes are waiting for the disk to service their requests.

Tuning the media server

This chapter includes the following topics:

- [Tips for backing up many small files](#)
- [Use backup-to-disk folders](#)
- [Use the Deduplication Option](#)
- [Tips for including and excluding selections](#)
- [Limit the backup frequency](#)
- [Run verify operations outside of the backup window](#)
- [Use disk staging](#)
- [Keep media servers upgraded](#)
- [Increase disk performance](#)

Tips for backing up many small files

Backup Exec may take longer to back up many small files than a single large file.

The following tips may improve performance when backing up many small files:

- On Windows, make sure that virus scans are turned off (turning off scans may double performance).
- Turn off or reduce logging.
The Backup Exec logging utilities have the potential to affect the performance of backup and recovery processing. Logging is usually enabled temporarily, to troubleshoot a Backup Exec problem. The amount of logging and its verbosity level can affect performance.
- Consider upgrading NIC drivers as new releases appear.

- When initially configuring the Windows server, optimize TCP/IP throughput as opposed to shared file access.
- Always boost background performance on Windows versus foreground performance.
- Regularly review the patch announcements for every server operating system. Install the patches that affect TCP/IP functions, such as the correction of out-of-sequence delivery of packets.

Use backup-to-disk folders

The following actions may improve performance when you use backup-to-disk folders:

- Set the maximum size for backup-to-disk files to an appropriate size. If you create small but numerous files, performance may slow since the computer must still process each file. However, if you create large files, file system limitations can cause memory allocation problems or network issues. These issues can be a problem if you store files across a network.
- Enable the option **Allocate the maximum size for backup-to-disk files** to reduce disk fragmentation. To ensure that backup data fills each backup-to-disk file to capacity, increase the append period of the media set that you associate with the backup-to-disk files. The backup data's overwrite protection period may also increase because the overwrite protection period starts at the end of the last append job.

Note: Backup performance may be affected when you enable the option **Allocate the maximum size for backup-to-disk files**. The job takes longer when Backup Exec creates a preallocated backup file at the start of a backup. Therefore, to gain the most benefit from preallocation, your backup strategy should include the reuse of these media by overwriting them. Then, the overhead of creating the preallocated backup file is limited to one time.

Additionally, when you enable the option **Allocate the maximum size for backup-to-disk files**, performance may increase if you disable the option **Buffered writes**. The combination of preallocation and buffered writes decreases performance on many systems.

To find what works best for your environment, enable the option **Allocate the maximum size for backup-to-disk files** for a first time backup job. Then, compare the throughput rate for that job with the throughput rate for the same job for which this option is not enabled.

- Experiment with the options for buffered reads and buffered writes on the media server. Enabling these options may increase backup performance.

Use the Deduplication Option

Instead of always running full backups, sending differential and/or incremental backups to a deduplication storage folder can improve performance. For full backups, Backup Exec compares each block with the deduplication database to determine if the data needs to be backed up. With a differential backup for example, Backup Exec can usually determine more quickly if a file needs to be backed up by checking the timestamp or the archive bit.

Tips for including and excluding selections

Consider the following when you create backup selection lists:

- Do not use excessive wild cards in file lists.
When wildcards are used, Backup Exec compares every file name against the wild cards, which decreases Backup Exec performance. Instead of placing `C:\Temp*` in an include or exclude list, use `/tmp/` or `C:\Temp`.
- Use exclude lists to exclude large, useless files.
Reduce the size of your backups by using exclude lists for the files that you do not need to preserve. For instance, you may decide to exclude temporary files. Use absolute paths for your exclude list entries, so that valuable files are not inadvertently excluded. Before adding files to the exclude list, confirm with the affected users that their files can be safely excluded. In a disaster, not being able to recover files costs much more than backing up extra data.

Limit the backup frequency

Minimize how often you back up the files that have not changed, and minimize your consumption of bandwidth, media, and other resources. To do so, limit full backups to monthly or quarterly, followed by weekly incremental backups and daily incremental backups.

Run verify operations outside of the backup window

You can set up a policy that contains a backup template and a verify backup sets template. The verify backup sets template verifies the integrity of the backup data after the backup is completed. You can schedule the verify operation to run at any convenient time outside of the backup window.

Running the verify operation outside of your backup window is helpful if your network resources are scarce. If you have difficulty completing backups within the allotted time window, you can schedule the verify operation to run at a more convenient time.

A verify backup sets template is especially beneficial if you use Backup Exec's Deduplication Option. If you use the default Verify after backup option, Backup Exec must perform the verify operation on both the server-side and the source-side. Backup Exec sends the backup sets across the network to verify them. The process can be time consuming depending on the amount of data that you are verifying and on your network configuration. If you schedule the operation using a verify backup sets template, the operation runs locally and more efficiently.

Use disk staging

With disk staging, also called a duplicate backup data job in Backup Exec, backups can be stored on disk initially, then copied later to another media type. The media type for the final destination is typically tape, but can be disk. This two-stage process leverages the advantages of disk-based backups in the near term, and preserves the advantages of tape-based backups for long term.

Symantec recommends that if you need to keep data longer than four weeks, create a duplicate backup job. A duplicate backup job copies the backup data from the original storage device to tape, which you can then send for long-term or off-site storage.

Disk staging can improve backup speed since there is not any tape device latency.

Keep media servers upgraded

To ensure the best possible performance, you should upgrade your media server to the latest Backup Exec release.

Upgrade the media server to avoid adversely affecting performance in the following areas:

- Catalog backup
- Reassignment of media to scratch media

Increase disk performance

You can use the following techniques to increase disk performance:

- Check the fragmentation level of the data.

A highly fragmented disk limits throughput levels. Use a disk maintenance utility to defragment the disk.

- Consider adding additional disks to the system to increase performance. If multiple processes attempt to log data simultaneously, divide the data among multiple physical disks.
- Determine if the data transfer involves a compressed disk. Windows drive compression adds overhead to disk read or write operations, with adverse effects on Backup Exec performance. Use Windows compression only if it is needed to avoid a disk full condition.
- Consider converting to a system with a Redundant Array of Inexpensive Disks (RAID). Though more expensive, RAID devices offer greater throughput and (depending on the RAID level) improved reliability.
- Determine what type of controller technology drives the disk. A different system might yield better results.

Tuning the Backup Exec data transfer path

This chapter includes the following topics:

- [About tuning the data transfer path](#)
- [Tuning suggestions for the data path](#)
- [Remote agent performance](#)
- [Network performance](#)
- [Media server performance](#)
- [About the Windows data buffer size](#)
- [Storage device performance](#)

About tuning the data transfer path

What limits the overall performance of Backup Exec is the slowest component in the backup environment. For example, a fast tape drive that is combined with an overloaded server yields poor performance. A fast tape drive on a slow network also yields poor performance.

The backup system is referred to as the data transfer path. The path usually starts at the data on the disk and ends with a backup copy on tape or disk.

The standard Backup Exec data transfer path consists of four components:

- The remote agent
- The network
- The media server

- The storage device

Ensure that your system meets Backup Exec's recommended minimum requirements. Refer to the *Symantec Backup Exec Administrator's Guide* for information about these requirements. Additionally, Symantec recommends that you have the most recent Backup Exec software patches installed.

Many performance issues can be traced to hardware or other environmental issues. You must understand the entire data transfer path to determine the maximum obtainable performance in your environment. Poor performance is often the result of poor planning, which results from unrealistic expectations of components of the transfer path

The requirements for database backups may not be the same as for file system backups. This information applies to file system backups unless otherwise noted.

If you want help fine-tuning a specific Backup Exec installation, please contact Symantec Consulting Services.

Tuning suggestions for the data path

In every backup environment, there is always room for improvement. To obtain the best performance from a backup infrastructure is not complex, but it requires careful review of the many factors that can affect processing. The first step is to gain an accurate assessment of each hardware, software, and networking component in the backup data path. Many performance problems are resolved before you change Backup Exec parameters. It is essential to thoroughly test both backup and restore processes after you make any changes to the Backup Exec configuration parameters.

Consider the following tips to improve backup environment performance and to avoid bottlenecks:

- Stripe a disk volume across drives.

A striped set of disks can pull data from all drives concurrently, to allow faster data transfers between disk drives and tape drives.

- Maximize the use of your backup windows.

You can configure all your incremental backups to happen at the same time every day. You can also stagger the execution of your full backups across multiple days. Large systems can be backed up over the weekend while smaller systems are spread over the week. You can start full backups earlier than the incremental backups. They might finish before the incremental backups and return all or most of your backup window to finish the incremental backups.

- Use local backup networks to decrease backup times and network traffic.

Specify a local backup network for operations between the media server and the remote computer to reduce backup time and reduce or eliminate network traffic on your enterprise networks. In addition, you can convert to faster technologies and back up your systems at any time without affecting the enterprise network's performance. This approach assumes that end users do not mind the system loads while backups take place.

- Avoid a concentration of servers on one network.
If many large data servers back up over the same network, install Backup Exec on some of them or attach them to private backup networks. Either approach decreases backup times and reduces network traffic for your other backups.
- Use dedicated backup servers to perform your backups.
For a media server, use a dedicated server for backups only. Using a server that also runs several applications unrelated to backups can severely affect the performance window and maintenance window.
- Consider the requirements of backing up your catalog.
Remember that the Backup Exec catalog needs to be backed up. To facilitate Backup Exec catalog recovery, the media server should have access to a dedicated tape drive, either stand-alone or within a robotic library.
- Try to level the backup load.
You can use multiple drives to reduce backup times. To spread the load across multiple drives, you may need to reconfigure the streams.

Remote agent performance

To improve the overall performance of Backup Exec, consider the components and factors in the table below.

Table 4-1 Remote agent performance

Item	Description
Disk fragmentation	Fragmentation severely affects the data transfer rate from the disk. Repair fragmentation by using disk management utility software.
Number of disks	Add disks to the system to increase performance. If multiple processes attempt to log data simultaneously, divide the data among multiple physical disks.

Table 4-1 Remote agent performance (*continued*)

Item	Description
Disk arrays	Convert to a system that is based on a Redundant Array of Inexpensive Disks (RAID). RAID devices generally offer greater throughput and (depending on the RAID level) improved reliability.
Controller technology	Consider if a different type of controller technology that is used to drive the disk can yield better results.
Backup Exec software location	If the data that is backed up is on the same physical disk as the Backup Exec installation, performance may be adversely affected, especially if Backup Exec debug log files are generated. If logs are used, the extent of the degradation is greatly influenced by the logs' verbose setting. If possible, install Backup Exec on a separate physical disk to avoid disk drive contention.
Hardware or software snapshots	If snapshots are taken before the backup of data, the time that is needed to take the snapshot can affect the performance.

Network performance

To improve the overall performance of Backup Exec, consider the network components and factors in the table below.

Table 4-2 Network performance

Item	Description
Network interface settings	<p>Use the following settings for the network interface:</p> <ul style="list-style-type: none"> ■ Set the network interface cards (NICs) to full-duplex. ■ Set both ends of each network cable (the NIC card and the switch) identically for speed and mode. Set both the NIC and the switch to full duplex, or errors may occur. ■ If auto-negotiate is used, set both ends of the connection to the same mode and speed. The higher the speed, the better. ■ Set all routers to full duplex. <p>Note: Using AUTOSENSE may cause network problems and performance issues.</p> <p>Consult the operating system documentation for instructions on how to determine and change the NIC settings.</p>
Network load	<p>To evaluate remote backup performance, consider the following:</p> <ul style="list-style-type: none"> ■ The amount of network traffic. ■ The amount of time that network traffic is high. <p>Small bursts of high network traffic for short durations can decrease data throughput rate. However, if the network traffic remains high, the network is probably the bottleneck. Try to schedule backups when network traffic is low. If your network is loaded, you may want to implement a secondary network which is dedicated to back up and restore traffic.</p>
Using multiple network interfaces	<p>If you have the Central Admin Server Option installed, you can distribute jobs over several network interfaces to improve performance. Managed media servers can use alternate network interfaces to run backup jobs.</p>

Media server performance

To improve the overall performance of Backup Exec, consider the components and factors in the table below.

Table 4-3 Media server performance

Item	Description
Hardware	<p>The following hardware conditions may cause slow throughput performance:</p> <ul style="list-style-type: none"> ■ The speed of the disk controller. ■ Errors that are caused by the disk drive, the tape drive, the disk controller, or the SCSI bus. ■ Improper cabling and termination. <p>Confirm that the controller is rated for the tape hardware and that the SCSI BIOS settings are set properly. Newer models of SCSI Controllers are set to communicate with SCSI hard drives by default.</p> <p>Use the following SCSI settings:</p> <ul style="list-style-type: none"> ■ Enable Disconnect and Initiate Sync Negotiation are set to No (in most cases). ■ Initiate Wide Negotiation is set to Yes when the tape device is connected to a 68-pin wide SCSI cable connector. <p>Ensure that tape drives are not connected to a SCSI Raid Controller.</p>
System	<p>The capacity and speed of the media server performing the backup, or the remote system being backed up significantly affects performance. System activity during backup also affects performance.</p> <p>Fragmented disks take a longer time to back up. Heavily fragmented hard disks affect the rate at which data is written to tape and also affect the overall system performance. Fragmented files take longer to back up because each segment of data is located at a different location on the disk. As a result, the disk takes longer to access the data. Make sure that you defragment disks on a regular basis.</p>

Table 4-3 Media server performance (*continued*)

Item	Description
Memory	<p>The amount of available memory affects backup speed. Insufficient memory, improper page file settings, and a lack of free hard disk space can cause excessive paging and slow performance.</p>
Compression	<p>Successful compression can increase the tape drive's data transfer rate up to twice the native rate. Compression can be highly variable depending on your input data. Compression algorithms look for the repeatable data patterns that can be compacted.</p> <p>Higher and lower compression occur depending on the type of files being backed up. Image and picture files are fully compressed on disks. Therefore, no hardware compression takes place during the backup, allowing the tape drive to operate at its native (non-compression) rate of speed. Image files from a graphical program like Microsoft Paint may compress at 4.5:1 or more, while binary files may compress at 1.5:1. Data that has already been compressed or data such as encrypted data or MPEG files may expand by about five percent if you attempt further compression. This expansion can reduce drive throughput.</p>
Files	<p>The total number of files on a disk and the relative size of each file affects backup performance. Fastest backups occur when the disk contains fewer large-size files. Slowest backups occur when the disk contains thousands of small files. A large number of files that are located in the same directory path back up more efficiently than backing them up from multiple directory locations.</p> <p>See “Tips for backing up many small files” on page 35.</p>

Table 4-3 Media server performance (*continued*)

Item	Description
Block Size	<p>Larger block sizes improve the compression ratio, which helps the drive to achieve better throughput and more tape capacity. Make sure that the block size and the buffer size are set properly. The throughput increases in proportion to the compression achieved, until the drive's maximum throughput is reached.</p> <p>Some devices such as LTO devices can provide better performance when larger block sizes are used. If you use a device that supports larger block sizes, you can change the device's block size in the Device Properties dialog box.</p> <p>For more information about using larger block sizes, see the following URL: http://seer.entsupport.symantec.com/docs/311736.htm</p> <p>You should also see the device manufacturer's documentation for help to configure the device.</p>
Network	<p>The speed of the physical connection limits the backup speed for a remote disk.</p> <p>The rate at which a remote server's hard disks can be backed up depends on the following:</p> <ul style="list-style-type: none"> ■ The make and model of network cards. ■ The network card driver. ■ The mode or the frame type configuration for the adapter. ■ The connectivity equipment (hubs, switches, routers, and so on). ■ The settings for Microsoft Windows. <p>Local disk drives on the media server can usually be backed up at a higher rate of speed than backing up remote servers across a network. Software compression can increase performance by reducing the amount of data that is transmitted across the network for a remote agent backup.</p> <p>See “Network performance” on page 44.</p>

About the Windows data buffer size

The size limit for data buffers on Windows is 1024 KB. This size is calculated as a multiple of operating system pages (one page = 4 KB). Therefore, the maximum is 256 OS pages, from 0 to 255 (the hex value 0xFF). A larger setting defaults to 64 KB, which is the default size for the scatter/gather list.

The setting of the maximum usable block size is dependent on the host bus adapter (HBA) miniport driver, not on the tape driver or the OS. For example, the readme for the QLogic QLA2200 card contains the following:

* MaximumSGList

Windows includes enhanced scatter/gather list support for very large SCSI I/O transfers. Windows supports up to 256 scatter/gather segments of 4096 bytes each, for transfers up to 1048576 bytes. The OEMSETUP.INF file has been updated to automatically update the registry to support 65 scatter/gather segments. Normally, no additional changes are necessary: This setting typically results in the best overall performance.

You can use the Backup Exec Discover tool to determine the maximum allowable block size for a tape drive.

You can find more information about the Discover tool at the following URL:

<http://seer.entsupport.symantec.com/docs/323012.htm>

Storage device performance

To improve the overall performance of Backup Exec, consider the components and factors in the table below.

Table 4-4 Storage device performance

Item	Description
Tape drive capacity	Tape drive wear is much less, and efficiency is greater, if the data stream matches the tape drive capacity and is sustained. Match the number of drives and the throughput per drive to the speed of the SCSI/FC connection, and follow the hardware vendors' recommendations.

Table 4-4 Storage device performance (*continued*)

Item	Description
Media positioning	When a backup or restore is performed, the storage device must position the tape so that the data is over the read and write head. The positioning can take a significant amount of time. When you conduct performance analysis with media that contains multiple images, allow for the time lag that occurs before the data transfer starts.
Tape streaming	If a tape device is used at its most efficient speed, it is "streaming" the data onto the tape. If a tape device is streaming, the media rarely has to stop and restart. Instead, the media constantly spins within the tape drive. If the tape device is not used at its most efficient speed, it may continually start and stop the media from spinning. This behavior is the opposite of tape streaming and usually results in a poor data throughput.

Table 4-4 Storage device performance (*continued*)

Item	Description
Data compression	<p>Most tape devices support some form of data compression within the tape device itself. Compressible data (such as text files) yields a higher data throughput rate than non-compressible data, if the tape device supports hardware data compression.</p> <p>Tape devices typically come with two performance rates:</p> <ul style="list-style-type: none">■ Maximum throughput Maximum throughput is based on how fast compressible data can be written to the tape drive when hardware compression is enabled in the drive.■ Nominal throughput Nominal throughput refers to rates achievable with non-compressible data. <p>Backup Exec cannot set the tape drive data compression ratio. Backup Exec can only enable and disable the compression functionality. Follow the instructions that are provided with your operating system and tape drive. In general, tape drive data compression is preferable to software compression. Software compression may be desirable for reducing the amount of data that is transmitted across the network for a remote agent backup.</p>

Additional Resources

This appendix includes the following topics:

- [Additional Resources](#)

Additional Resources

More information about Backup Exec tuning and performance is available from the following sources:

- Checking system performance by using Service Performance Monitor:
<http://www.symantec.com/docs/TECH124254>
- Using Rsync to test disk I/O performance issues when using the Remote Agent for Macintosh Servers:
<http://www.symantec.com/docs/TECH87857>
- Correcting slow backup performance and agent initialization problems on fragmented Windows Server hard disk partitions:
<http://www.symantec.com/docs/TECH11562>

