WHITE PAPER: DATA PROTECTION

Best Practices for Using Symantec™ Online Storage for Backup Exec

Confidence in a connected world.

symantec™

White Paper: Data Protection

Best Practices for Using Symantec™ Online Storage for Backup Exec

Contents

Executive summary

Symantec Backup Exec™ for Windows® Servers is a comprehensive backup and recovery solution that protects a wide range of data resources, from file systems to popular databases. Symantec Online Storage enhances Backup Exec by providing secure off-site backup storage hosted in redundant data centers that are managed by Symantec for long-term storage and disaster recovery purposes. Symantec Online Storage for Backup Exec is an offering of Symantec Protection Network (**www.spn.com**), which provides market-leading Symantec technologies to customers as online services.

Symantec Online Storage for Backup Exec is completely integrated with Backup Exec for Windows Servers (v12 and higher); in theory, you could send any on-premise backup to Symantec Protection Network data centers. However, when backing up over the Internet to managed data centers, the performance is not the same as when backing up to local disk. The amount of bandwidth available, frequency of backups, and amount of data transferred will play a significant role in performance using Symantec Online Storage. In addition, costs for the service are based on total storage allocated and consumed, and storage can accumulate quickly if not managed carefully.

This paper suggests best practices to consider when planning and implementing off-site backups using Symantec Online Storage for Backup Exec. These suggestions are designed to help you achieve your data protection objectives while balancing bandwidth and cost considerations.

Symantec Online Storage for Backup Exec

Symantec Online Storage extends the key capabilities of Symantec Backup Exec (v12 and higher), adding the ability to store copies of backups in secure, redundant data centers managed by Symantec. Online Storage gives businesses a simple, secure, and cost-effective option for storing backup data off-site without the inconvenience and cost of tape vaulting services or other forms of off-site tape storage.

Symantec Online Storage is tightly integrated with Symantec Backup Exec. To understand its role, it helps to understand the existing Backup Exec components with which it interacts.

Backup Exec basics

Every Backup Exec implementation will have one or more Media Servers, which are the “brains” of a Backup Exec deployment. The Media Server performs job configuration and scheduling, and controls data movement to and from agents and storage devices/media.

Other important terms and components include:

- **Agents**—Agents are software components installed on resources that need to be protected, including file servers, database servers, and application servers.
- **Storage devices**—Backup Exec can write to and retrieve data from either tape drives (standalone or in libraries) or disk devices.

- **Storage media**—When data is written to the storage device, it is captured on physical tapes (when using tape devices), CDs, DVDs, or disks (when using hard drives).
- **Backup to Disk**—Disk-based backups use Backup to Disk (B2D) folders. The files written to the B2D folder are often referred to as "virtual tapes." Virtual tapes are typically 1 GB in size and carry the extension ".BKF."

Symantec Online Storage works within this Backup Exec environment, supplementing on-premise backups with backups sent over the Internet to Symantec Protection Network data centers.

Using Online Storage with Backup Exec

To use Symantec Online Storage, you must connect to the Symantec Protection Network portal, create an account, and select a service plan that meets your off-site storage requirements. You also need to add two components to the backup environment:

- **Symantec Protection Agent**—When installed on the Backup Exec Media Server, this small agent acts like a Backup Exec storage device that writes to the Symantec Protection Network.
- **Storage Device for Online Storage**—Once you have installed the Symantec Protection Agent, you can create and define a backup folder that looks and functions like a Backup to Disk (B2D) folder, except that the files (virtual tapes) created are transported over the Internet and stored, encrypted, in secure and redundant Symantec Protection Network data centers.

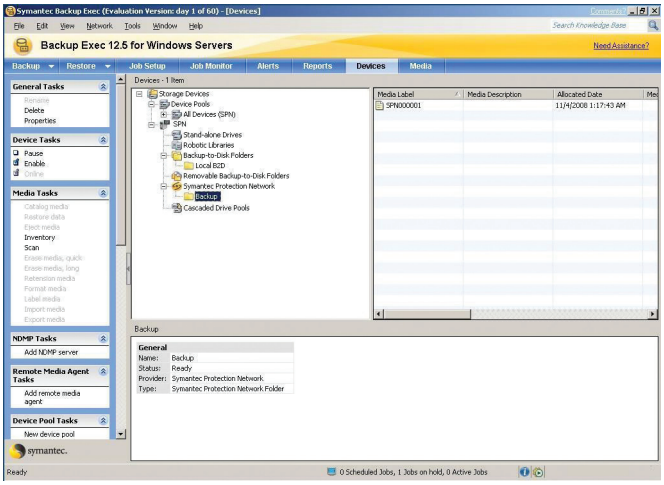

Figure 1. Symantec Protection Network appears on the Devices list

Backup Exec uses the Duplicate Backup Set function to send backups to Online Storage. The Duplicate Backup Set function creates a byte-for-byte copy of an existing physical or virtual tape and sends it over the Internet to the Symantec data centers hosting the Online Storage service.

Key considerations

Several factors may influence your decisions about how and when to use Symantec Online Storage. These include the following:

- **Duplication**—The Online Storage backup must be a duplicate of an existing backup set. For optimal performance, the backup set being duplicated should be a local disk backup (B2D folders). It is possible to duplicate a tape backup for Online Storage, but using tape will slow the backup performance.
- **Cost**—Cost of the service is based on total storage consumed, with pricing by storage tiers. You can upgrade storage tiers as your needs grow. You can control the cost of the service by limiting the amount of data you store in Symantec's managed data centers.
- **Bandwidth**—Bandwidth is a key limitation. Backups to Online Storage take place over the available Internet connection, and bandwidth affects overall backup and recovery performance. For example, a T3 line offers significantly more bandwidth than a T1 line or high-speed cable modem. Very large backups may affect the Internet bandwidth available for other services as well.

Bandwidth and storage size are significant considerations when it comes to selecting the size and frequency of backups to the Online Storage Service. For example, consider a typical customer with 100 gigabytes (100 GB) of data that is protected on a regular basis with Backup Exec. Most customers will use some variation of a GFS (Grandfather, Father, Son) backup rotation that looks something like the following:

- Take a full backup at the beginning of each week
- Perform incremental backups during the remaining days of the week
- At the end of the month, keep the last full backup, but erase or repurpose the rest of the tapes (or virtual tapes) to be used for the next month's rotation

The GFS rotation allows for a relatively rich backup history, but generates a significant amount of data to be stored on tape or disk. The table below summarizes the amount of data that may be generated in a typical 100 GB scenario over the course of a month.

Table 1. Data generated by GFS rotation with 100 GB data set

Backup Type	Base Data Size	Incremental Data Size (5 GB change each day)	Total Data on Storage Media
Week 1			
Full (Sunday)	100 GB	0	100 GB
Incremental (Mon–Sat)	0	35 GB	135 GB
Week 2			
Full (Sunday)	100 GB	0	235 GB
Incremental (Mon–Sat)	0	35 GB	270 GB
Weeks 3 & 4			
Full (x2)	200 GB	0	470 GB
Incremental (x12)	0	60 GB	530 GB

The second month will start over again, with the addition of the 100 GB from the previous month's full backup. It is evident that, when this strategy is used, storage accumulates quickly. While this strategy works well for on-site backups, it can overwhelm available Internet bandwidth and will incur increasing costs. Consequently, it is best to limit the size and frequency of the backups that are copied to off-site storage.

Best practices for Symantec Online Storage

The practices suggested below are designed to help you balance bandwidth and cost considerations against your need for long-term data storage and off-site storage for disaster recovery purposes.

Complement existing, on-premise backups

Symantec Online Storage is designed to be used for long-term storage and disaster recovery purposes, not as primary backup storage. While you can use Symantec Online Storage for any kind of backup, recovering data from local disk is always going to be much faster than recovering from off-site data centers over the Internet. Symantec Online Storage backups should supplement—not replace—your on-premise backups.

Take the time to create a backup strategy that uses Online Storage resources as part of your overall backup plan, balancing cost and performance with the need for long-term protection and off-site storage. Once you have identified your strategy, putting it into action takes only minutes.

Identify your most essential data for Online Storage

Because it is so easy to send backups to the Symantec Protection Network, you might be tempted simply to send everything off-site.

In practice, however, we suggest that you use Online Storage for your most essential data: data that needs long-term storage, or data that simply must be available in case of a site-wise disaster. Small, critical backup sets are the best candidates. These may include:

- Essential databases (customer relationship management databases)
- Customer records or regulated data
- Employee or payroll information
- Backup Exec catalogs

If nothing else, you should be sure to exclude the files that you do *not* need stored in secure, redundant data centers, including many users' desktop files and MP3 files.

Protect the Backup Exec catalogs

By sending your Backup Exec catalogs to Symantec's secure data centers, you can protect your data from the loss of the Backup Exec Media Server. If the Media Server is lost, you can install Backup Exec and the Symantec Protection Agent on a new server, and then restore the Backup Exec catalogs from the Symantec data centers. You will need the encryption passphrase to recover the data.

Create a distinct backup set to duplicate for Online Storage

The Symantec Online Storage backup set must be a duplicate of another backup set. Simply duplicating an existing set is the easiest thing to do, and in some cases may be the right thing. But we suggest that you look carefully at your needs and consider defining distinct backup sets for Online Storage backups.

There are several benefits to having a distinct backup set that you duplicate for Symantec Online Storage backups. You can:

- Simplify management and avoid impacting your current setup
- Select only your most mission critical data for off-site backup
- Monitor storage usage and alter the frequency or size of the Symantec Online Storage backups, without affecting on-premise backup procedures, allowing you to better manage costs and retain flexibility in your environment
- Erase specific backups from Symantec Online Storage media on premise and from your off-site storage, again without affecting the core backups you maintain locally.

Set different policies for online storage retention

You can retain backup data in Symantec Protection Network data centers for as long as you need it.

Once you have created the duplicate backup set that you send to Online Storage media, you should set distinct retention policies for this set, based on your business and data needs. These policies will allow you to manage your storage consumption and subscription fees.

Long-term, protected storage is one of the key benefits of Symantec Online Storage. It provides an easy-to-manage alternative to tape. And because it reduces or even eliminates the cost of backup-related vaulting services, media, and hardware, it can be a lower-cost solution.

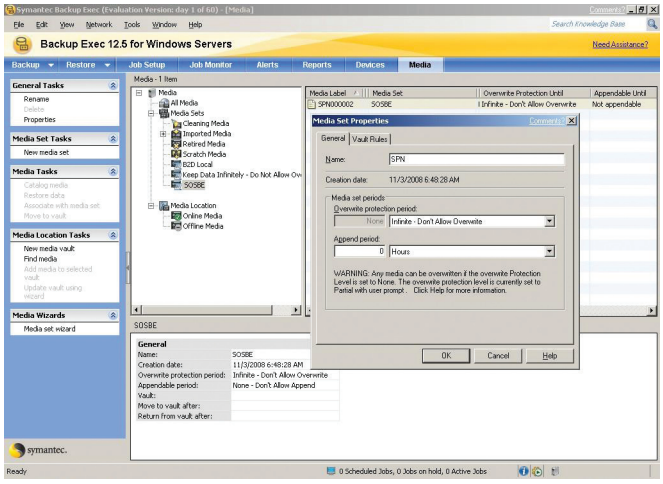

Figure 2. Assigning Media Properties to the "SPN" device

Use incremental and differential backups with care

Incremental and differential backups are one way to reduce the amount of data being backed up.

- **Incremental backups** write only the data changed since the last backup of any kind.
- **Differential backups** write the data changed since the last full backup.

The distinction is an important one when it comes to recovery. If you perform a full backup on Monday and incremental backups each day of the week, then to recover from a failure on Friday you need to restore one full backup and four incremental ones. Given the same failure, if you take differential backups each day, you would need to retrieve only the full backup plus Friday's differential.

When using either incremental or differential backups with Symantec Online Storage, keep the following in mind:

- You cannot start a new incremental or differential backup until the prior Symantec Online Storage backup job has completed. For very large backups over relatively low-bandwidth connections, this can be an issue.
- Recovery from incremental backups can take longer because of the larger amount of data you must retrieve—and the inherent delays of Internet access are, again, a potential bottleneck. Be sure to intersperse incremental backups with full backups.
- If you use Online Storage primarily for long-term storage and disaster recovery capabilities, then less frequent full backups may make more sense than more frequent incremental backups.

Track your storage usage

With the Symantec Online Storage service, you pay for total storage consumed, so tracking your storage consumption is a good idea. You can track the amount of storage you are consuming with the service from the Symantec Protection Network portal.

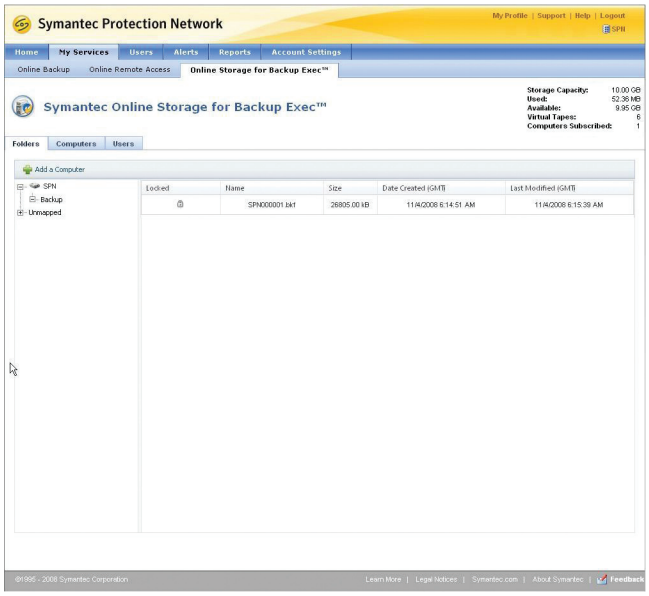

Figure 3. Tracking storage usage from the Symantec Protection Network portal

If your storage usage is growing faster than expected, either restrict the volume of data you send off-site or increase your storage allowance on the service. It is more cost-effective, in general, to increase your storage allowance than to pay overage charges for exceeding the storage volume allotted to your service level.

You can also control storage growth by deleting the files you have stored off-site from within Backup Exec.

Erase unneeded files to reclaim space

If you no longer need backup sets stored off-site, you can erase them from within Backup Exec. From the list of available devices, simply select the Symantec Protection Network folder containing the data that you want to erase, and then pick the Erase media task. This will erase the media stored in the Symantec data centers and reduce the amount of your allocated storage that you are using.

Protect the passphrase!

All backups to Symantec Protection Network data centers are transmitted and stored in an encrypted format, using the encryption capabilities of Backup Exec. Even the Symantec staff cannot access the encrypted data, so the privacy of your data is protected.

You provide an encryption passphrase when creating the backup set. The Backup Exec Media Server maintains that passphrase. As long as you are restoring backups using that same Media Server, you do not need to provide the passphrase on recovery.

However, if you need to restore data to a different system (for example, if the Media Server itself was lost), then you must provide the passphrase to decrypt the data. You should create—and follow—policies for protecting the passphrase:

- Make sure that the passphrase is written down and stored securely somewhere.
- Ensure that the passphrase will be accessible in the case of a site-wide disaster.
- Carefully monitor and control who has access to the passphrase. More than one person must be able to access it.

Summary

With Symantec Protection Network, Symantec makes enterprise-class services available to businesses of all types and sizes. Symantec Online Storage for Backup Exec gives your business immediate and easy access to managed storage in secure, redundant data centers.

Symantec Online Storage for Backup Exec is best used as part of a comprehensive data protection strategy that includes on-premise backups for localized failures as well as long-term, off-site storage for disaster recovery and long-term retention. The best practices highlighted in this paper should help you fit the off-site storage capabilities of Symantec Protection Network into your Backup Exec environment and processes.

About Symantec

Symantec is a global leader in providing security, storage, and systems management solutions to help businesses and consumers secure and manage their information. Headquartered in Cupertino, Calif., Symantec has operations in more than 40 countries. More information is available at: www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

12/08 14173820