

Data Protection for Microsoft SharePoint Portal Server 2003

VERITAS Backup Exec™ 10 for Windows Servers

Agent for Microsoft SharePoint Portal Server

Applicable for Windows SharePoint Portal Server Environments
For Use with Backup Exec 10 for Windows Servers and Microsoft
Windows Server 2003

TABLE OF CONTENTS

Executive Summary	3
Product Highlights	3
How it works	4
Conceptual Overview	4
Data-Protection Planning	5
Backup Exec 10 Configuration.....	5
Backup Exec 10 Media Server	5
Backup Exec 10 Remote Agent for Windows Servers	5
Backup Exec 10 Media Sets.....	5
Configuring Intelligent Disaster Recovery	5
Monitoring and Reporting	6
Backing Up And Restoring SharePoint Configuration Databases	8
Backing Up The Portal Site Databases.....	8
Single Sign-On Database	9
SharePoint Portal Server 2001 Agent.....	9
Restoring SharePoint Resources	10
Restoring To The Same Database Server.....	10
Performing A Redirected Restore	10
Limitations on Restore	10
Disaster Recovery	10
Single Sign-on Services Database	12
Summary	12

EXECUTIVE SUMMARY

Microsoft SharePoint Portal Server (SPS) 2003 offers a second-generation document management, project collaboration and intranet site management tool with improved scalability and flexibility for Windows servers. It facilitates easy organization, sharing, retrieval and publishing of information over corporate intranets and seamlessly integrates with Microsoft Office and Web-development tools.

The release of SharePoint Portal Server 2003 and Windows SharePoint Services has dramatically changed the product's architecture, requiring new methods for presenting the data to the user as well as for backing up this data. This format is not only different from the previous version of Microsoft SharePoint technologies but also introduces many complexities in supporting data that is passed freely between servers. The exchange of information between servers in the farm will not only change how you back up this data but how the user goes about selecting backup and restore methods without having to know the details about the server farms' configuration.

This white paper provides a detailed review of how to fully protect SharePoint Portal Server 2003 on the Windows Server 2003 platform using Backup Exec 10 *for Windows Servers* version and the optional VERITAS Backup Exec Agent for Microsoft SharePoint Portal Server (SPS agent).

PRODUCT HIGHLIGHTS

The new Backup Exec SPS agent automates the many steps required to fully protect your SPS environment. The methods involved in protecting SharePoint Portal Server 2003 require the use of the Backup Exec Media Server and the SPS agent.

The SPS agent supports backup and restore of SharePoint related databases, optional document libraries and some additional metadata. The SPS 2003 agent lets administrators scale a single-server SharePoint Portal Environment to large server farm environments. The use of server farms lets an administrator distribute the various components of a SPS 2003 configuration across many servers in an enterprise. The SPS agent then depicts a concise view of the server farm topology independent of the rest of the servers in the enterprise.

KEY BENEFITS

- Supports Microsoft SharePoint Portal Server 2003 and Serer Farm
- Fully automated
- Simplifies and ensures that all the necessary components of SharePoint Portal Server 2003 are fully protected and available for restoration and complete disaster recovery

HOW IT WORKS

CONCEPTUAL OVERVIEW

A SharePoint Portal Server 2003 (SPS 2003) deployment contains one or more servers configured in either single server or server farm topologies. SPS 2003 consists of the following minimum components: one instance of Microsoft MSDE or SQL server, a Web/search server, an index/job database, and optional backward compatible document libraries.

To clarify various SPS 2003 configurations, note these standard deployments for establishing common terminology:

Standard SPS 2003 deployments

- **Single server:** Single server hosting all SharePoint components, including either MSDE or SQL database.
- **Small server farm:** All SPS components on a single server except for the SQL database, which has been installed on a separate server
- **Medium server farm:** One or more front-end Web servers with the search component enabled. One or more index servers, one of which is the job server. One or more servers running SQL 2000.
- **Large server farm:** Two or more front-end Web servers. Two or more search servers. One or more index servers, one of which is the job server. One or more servers running SQL 2000.

Note: In all types of deployments, if SQL is deployed, the optional backward compatible document library can be hosted on a separate server.

DATA-PROTECTION PLANNING

Without the proper tools and processes in place, the time-consuming tasks of collaborating, publishing, and controlling access to documents within an organization could easily result in data being lost, overwritten, duplicated, or misplaced. While SharePoint Portal Server 2003 solves these and other problems, it does not use adequate data-protection tools for reliable disaster recovery, scalability, and ease of use. Without the proper data protection strategy, an organization places its documents and data at risk — the environment has no defined data-protection schemes, and recovery processes have not been defined. It is crucial that organizations research, evaluate, and deploy a complete data-protection solution for SharePoint Portal Server 2003.

As organizations deploy SharePoint Portal Server 2003, the common question of “How to effectively protect valuable data stored within SharePoint Portal Server 2003?” will arise. This white paper answers this simple, yet crucial question. It also presents various tools, processes, and strategies available to back up and restore SharePoint Portal 2003 servers. Each organization must decide — based on size, infrastructure, and type of SharePoint Portal Server 2003 deployment — what combination of these tools best fits its environment. Additionally, when determining specific SharePoint Portal Server 2003 data protection needs, organizations must consider these questions:

- Will backup processes be performed while SharePoint Portal Server 2003 is on-line or off-line?
- Will backup processes be performed from a central location or distributed among multiple servers?
- Will backups be stored on tape media or disk volumes?
- Will backup processes be performed individually or combined with other protected resource backups such as Exchange, SQL, or Domino?
- How frequently should a backup of SharePoint Portal Server 2003 be performed?
- How can a corrupted or accidentally deleted files or databases be recovered?
- How is protecting a SharePoint Portal 2003 server farm different from a single server installation?
- What tools exist to help automate and simplify SharePoint Portal Server 2003 data protection?
- What steps are involved for quickly and reliably recovering from catastrophic data loss?

BACKUP EXEC 10 CONFIGURATION

Backup Exec 10 Media Server

Before you can begin implementing your SharePoint Portal Server 2003 data-protection plan, you must have installed and configured a Backup Exec 10 media server. Installation and configuration instructions are available in the VERITAS Backup Exec 10 for Windows Servers Administrator's Guide. The media server must have the following options licensed:

- Intelligent Disaster Recovery (if part of plan)
- Backup Exec Agent for Microsoft SharePoint Portal Server

Backup Exec 10 Remote Agent for Windows Servers

Before you can create jobs, you must install the Backup Exec 10 Remote Agent for Windows Server on each Server that participates in the farm topology that will be backed up.

Backup Exec 10 Media Sets

After you have a media server available, depending on your media management scheme, you may want to define one or more media sets for the SharePoint Portal Server 2003 servers at your site. The media set controls the overwrite protection period, which is how long that data is retained before being eligible to be overwritten, and the append period, which is how long that data can be appended to media. Defining a media set lets you customize backup-job retention period.

Configuring Intelligent Disaster Recovery

If your data-protection plan includes the Intelligent Disaster Recovery (IDR) option for any of the servers, make sure to create an IDR bootable recovery image for each one. Please refer to the VERITAS Backup Exec 10 for

Windows Servers Administrator's Guide for detailed instructions on how to configure and maintain IDR recovery data.

Monitoring and Reporting

Regular monitoring of backup jobs is an important task for backup administrators. If backups of a SharePoint Portal Server 2003 server fail for any reason, it may not be possible to restore that server to its most recent state. For this reason, the SharePoint Portal Server administrator should also monitor the backup status.

Backup Exec 10 offers a management pack for use with the Microsoft Operations Manager (MOM) at no additional cost. The MOM management pack monitors the health and availability of Backup Exec for Windows Servers software. By detecting, alerting about, and automatically responding to critical conditions, the management pack helps identify, correct, and prevent possible service outages. The management pack is available for download at <http://seer.support.veritas.com/docs/272197.htm>.

Backup Exec 10 also includes more than 40 reports that show detailed information about protected servers, media, and devices. When generating most of the reports, you can specify settings that serve as filter parameters or a time range for the data that you want to include in the report. This makes it possible to create a report that includes the set of SharePoint Portal Server 2003 servers. You can run and view a new report immediately, or you can create a job that saves the report data in the job history. You can also view general properties for each report.

Backup Exec 10 can schedule a report to run at a specified time or on a recurring schedule, and it can distribute reports through email notifications. This makes it possible to run scheduled reports that supply the data-protection status of a set of SharePoint Portal Server 2003 servers, as well to distribute the reports to all members of the organization responsible for the maintenance of these servers.

Reports are generated using Crystal Reports and can be viewed and printed in an HTML file format. If Backup Exec detects that Adobe® Reader is available, it displays reports in the Adobe Portable Document Format (PDF). The free Adobe Reader software is available at <http://www.adobe.com/acrobat>.

Below is a list of Backup Exec 10 reports that will help backup and SharePoint Portal Server administrators effectively monitor the data-protection status of a set of SharePoint Portal Server 2003 servers:

- Backup Job Success Rate – shows the success rate of success for jobs run on a set of selected servers.
- Backup Resource Success Rate – shows the success rate of success for each resource on set of selected servers.
- Backup Set Details by Resource – shows detailed information for each resource backed up on a set of selected servers.
- Backup Sets by Media Set – shows detailed information about all backup sets on selected media sets.
- Failed Backup Jobs – lists failed jobs for a set of selected servers, over a user-definable time period.
- Media Set – lists all the media used for a user-selectable group of media sets.
- Overnight Summary – provides an easy-to-view list of all backups within the last 24 hours for a set of selected servers.
- Policy Jobs by Resource Summary – shows details about each resource backed up in a user-defined period using policy-defined jobs, for a set of selected servers.
- Policy Jobs Summary – shows all the jobs derived from selected policies in a specified time range.
- Policy Protected Resources – shows a list of resources, and the policy and templates assigned to them, for a set of selected servers.
- Problem Files – shows a list of files that Backup Exec had a problem backing up, by resource, for a set of selected servers.
- Resource Backup Policy Performance – shows the success rate of policy-derived jobs for a user-defined time period, on a set of selected servers.
- Resource Risk Assessment – provides a list of resources for which the most recent backup failed, for a set of selected servers.

- Restore Set Details by Resource – shows detailed restore information by resource, in a user-defined time period, and for a set of selected servers.

To view or schedule reports, open the Backup Exec 10 administration application, and click the **Reports** tab. Right-click a listed report to see its run and scheduling options.


LICENSING OF BACKUP EXEC IN THE SPS ENVIRONMENT

The SPS agent consists of the following components: One component to protect a single SQL database instance and one Remote Agent for Windows Servers (CAL).

In its simplest form, SPS 2003 is installed on a single server. The SPS server consists of the SharePoint application and an MSDE or SQL 2000 instance running on Windows Server 2003. The Backup Exec SPS agent includes all the necessary components to protect this configuration.

A typical small server farm consists of two servers: a server running SPS 2003 on Windows Server 2003 and a second server running SQL 2000 on Windows 2000 or Windows 2003. A single Backup Exec SPS agent license and a remote agent for Windows servers (CAL) are required to protect this farm configuration.

A medium or large server farm consists of a minimum of the following components: one or more SQL servers running on Windows server 2000 or 2003, one or more Web servers, one or more index/job servers, and one or more search servers. To fully protect these large server farm configurations, you need one SPS 2003 agent, an SQL server agent for each SQL server above one, and a remote agent for each additional server beyond the initial server in the farm.



In the example of a medium server farm above, the following licenses are required:

- One SPS server agent
- One additional SQL database agent for the second SQL database instance
- Five additional Remote Agents *for Windows Servers*

Note: In cases where you have multiple SPS farms sharing a single SQL server, each farm requires a separate Backup Exec SPS agent license.

BACKING UP THE COMPLETE SHAREPOINT PORTAL SERVER ENVIRONMENT

The best place to start your backup of the SPS environment is by pointing to the Farms node in the Remote Selections list. The Farms node shows the most complete view of the topology. You can automatically select the entire server farm to back up the entire set of servers and components for your environment. If you browse to local nodes, you may only see a partial view of the farm typology depending from which server you are viewing the configuration. Web servers show all resources in the local node but some may not be selectable. If the SharePoint configuration includes remote selections, a new Microsoft SharePoint Server Farms node will be added to the remote selections node after you install a Remote Agent for Windows Servers (CAL). You can add these nodes manually or automatically. To add these nodes manually, select **Add Server Farms** in the context menu. To add nodes automatically, browse to a front-end Web server that participates in a server farm.

BACKING UP AND RESTORING SHAREPOINT CONFIGURATION DATABASES

In Microsoft Office SharePoint Portal Server 2003, all server and site configuration information is stored in the configuration database, and all site content is stored in content database(s). If you want to individually restore all the Microsoft Office SharePoint Portal Server 2003 information on your server or server farm, you must back up these databases with the Backup Exec Agent for SPS 2003 Server as part of a full backup. If you have a server farm configuration, the Backup Exec SharePoint Portal agent contacts the front-end Web server. The Web server is running a process that queries all the database(s) (and/or servers) and collects the necessary data for backup. Since the front-end Web servers collect information from all the SharePoint Servers in the farm configuration, best practices dictate that administrators leverage the Intelligent Disaster Recovery (IDR) option on the Web servers to be able to rebuild these servers in the event of a catastrophic failure of the Web servers.

The configuration database name is determined after the Microsoft Office SharePoint Portal Server 2003 is installed; the default is SPS01_Config_db. For more information, see your system administrator or SharePoint documentation.

BACKING UP THE PORTAL SITE DATABASES

The portal site databases for Microsoft Office SharePoint Portal Server 2003 are created in the SQL instance or instances designated for these components in the farm topology. The database names consist of the first eight characters of the portal name with no spaces, an increment value, and concatenated with _PROF, _SITE, or _SERV. For example, the "Team Portal" portal site will consist of the following databases:

- TeamPort1_PROF
- TeamPort1_SITE
- TeamPort1_SERV

Best practices strongly encourage all administrators to backup all portal site databases, including index databases and any team databases at the same time.

- If you backup the configuration database, you must ensure that the profile, site, and server databases were all backed up at the same time. If all databases are restored, the Microsoft Office SharePoint Portal Server 2003 server or farm will be restored to same state that it was in when it was backed up. (This is due to the restore of the configuration database.) However, problems may result if the configuration database is restored individually without the other databases. If configuration information is lost, the Microsoft Office SharePoint Portal Server 2003 farm or server may be compromised, which may result in data loss. To ensure that the Microsoft Office SharePoint Portal Server 2003 server or farm can be restored in its entirety, you must include the configuration database when you back up all the databases. For example, the following databases must be backed up to ensure a complete restore of the configuration database.

- TeamPort1_PROF

- TeamPort1_SITE
 - TeamPort1_SERV
 - SPS01_Config_db
- If each of these components is on a separate server, each of these components will have its own backup set. Best practices for restore would be to select all sets in one job and bring all databases on-line after the restore job is completed for all databases rather than bringing each database on-line after the individual database restore is complete. See the Backup Exec administrators guide for details on how to configure your database restore to not automatically start up on completion of the restore.

IMPORTANT: Proceed with caution. Restoring these individual databases will only be useful if the configuration on the farm has not changed, such as server names, between the time of the backup and the restore. The configuration database holds all of the topology information. If you restore the configuration database and the topology changed then it will not be valid anymore. In that case you would be better off creating a new configuration database and restructuring the topology.

SINGLE SIGN-ON DATABASE

If the SharePoint Farm uses the Microsoft Single Sign-On Service, that database as well as the encryption key must be backed up. The encryption key is automatically included with the Single Sign-on database when it is backed up using the SPS agent.

SHAREPOINT PORTAL SERVER 2001 AGENT

The new SPS agent can be used to protect instances of SPS 2001 in conjunction with SPS 2003. Protecting SPS 2001 does not require the use of a separate agent.

RESTORING SHAREPOINT RESOURCES

You can restore the SharePoint resources to the same server from which they were backed up or you can redirect the restore to another location.

RESTORING TO THE SAME DATABASE SERVER

The SPS agent enables you to restore at many different levels. Administrators are able to restore at the farm level or at a portal site level. Since you can have many portal sites within a farm, it will be most common to restore at a portal site level to restore deleted or corrupted documents in the SPS 2003 environment. Running regular backups of your servers and sites lets you restore them in case there is a failure. Performing restores is covered in detail in the *Backup Exec for Windows Servers* administrator's guide. Please refer to this document for detailed procedures on restore.

PERFORMING A REDIRECTED RESTORE

Depending on the role of the server, you must ensure that the server detects that a restore of the data was performed and the local cache on the server is refreshed. For more information about adding individual servers to the farm, see the *Backup Exec for Windows Servers* administrator's guide.

LIMITATIONS ON RESTORE

- SPS 2003 does not support redirected restore of the config DB or the SSO DB
- Administrators do not have the ability to redirect resources from more than one portal site or team site at a time
- Administrators cannot redirect a portal to another portal in the same farm if the original portal site still exists in the farm
- When redirecting a portal site restore the target portal site must already exist and have the same structure. The redirected restore process replaces the data in the portal site but does not automatically create the necessary structure.
- Best Practices recommend that you do not restore the subcomponents of a portal site independently as there are commonly interdependencies between the subcomponents of a portal site.

DISASTER RECOVERY

A prerequisite to performing the steps below is to have created IDR media for disaster recovery. Recovery scenarios include general steps and there may be unique environments that require additional steps, see the *Backup Exec* administrators guide for details. The recovery of an MSDE or Microsoft SQL Server 2000 database that has become corrupted, the recovery or replacement of a single hard drive or server, or recovery of multiple servers following a site disaster as all covered below at a high level. Recovery of a database or single server is straightforward, but recovery following a site-wide disaster is more complex.

Stand-alone (MSDE) configuration

1. Utilizing the Intelligent Disaster Recovery media that was previously created, perform a recovery of the server (requires reboot).
2. From Backup Exec, restore the non-SPS SQL databases (Master, Model, MSDB, Pubs, etc.).
3. Disconnect from the SharePoint configuration database (required as part of the disaster recovery process).
 1. For MSDE based installations, you first need to change the "StandAlone" registry value by performing the following steps:
 1. Open RegEdit and navigate to:
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\SharePoint Portal Server].
 2. Change the value called "StandAlone" from "1" to "0".
 2. Disconnect from the SharePoint configuration database.

3. Change the "StandAlone" registry value back to "1".
4. Delete all the databases that are marked as 'suspect' within SQL.
5. Delete the Index/Search GUID folders (located under the SPS installation path).
6. Remove SharePoint Services from the virtual server(s).
7. Recreate the configuration database with the same name as before and in the same location.
8. Configure the SharePoint topology as it was before.
9. From Backup Exec, select and restore all the SharePoint site databases.
10. From Backup Exec, select and restore the Single Sign-on database (if applicable).
11. From Backup Exec, select and restore the Document Library Store (if applicable).

Single Server with SQL Server configuration

1. Utilizing the Intelligent Disaster Recovery media that was previously created, perform a recovery of the server (requires reboot).
2. From Backup Exec, restore the non-SPS SQL databases (Master, Model, MSDB, Pubs, etc.).
3. Disconnect from the SharePoint configuration database (required as part of the disaster recovery process).
4. Delete all the databases that are marked as 'suspect' within SQL.
5. Delete the Index/Search GUID folders (located under the SPS installation path).
6. Remove SharePoint Services from the virtual server(s).
7. Recreate the configuration database with the same name as before and in the same location.
8. Configure the SharePoint topology as it was before.
 - Note: If separate SQL instances (different than one used by configuration database) were used for the Content and/or Component Settings databases, then configure SharePoint to point to the other instance(s).
9. From Backup Exec, select and restore all the SharePoint site databases.
10. From Backup Exec, select and restore the Single Sign-on database (if applicable).
11. From Backup Exec, select and restore the Document Library Store (if applicable).

Individual Servers of a SharePoint Portal Server Farm (Small, Medium or Large)

Search Server (Large)

1. Utilizing the Intelligent Disaster Recovery media that was previously created, perform a recovery of the server (requires reboot).
2. Since nothing specific to SharePoint Portal Server is backed up from the Search server(s) in this case, no further action is required. However, you may need to restart the PS (Portal Server) Search service.

Index Server/Job Server (Medium or Large)

1. Utilizing the Intelligent Disaster Recovery media that was previously created, perform a recovery of the server (requires reboot).
2. Delete the Index GUID folders (located under the SPS installation path).
3. From Backup Exec, select and restore all Index databases associated with this Index server.

Web Server (Medium or Large)

1. Utilizing the Intelligent Disaster Recovery media that was previously created, perform a recovery of the server (requires reboot).
2. Since nothing specific to SharePoint Portal Server is backed up from the Web server(s) in this case, no further action is required.

Doc Library Server (All)

1. Utilizing the Intelligent Disaster Recovery media that was previously created, perform a recovery of the server (requires reboot).
2. From Backup Exec, select and restore the Document Library Store.

SQL Server (All)

1. Utilizing the Intelligent Disaster Recovery media that was previously created, perform a recovery of the server (requires reboot).
2. From Backup Exec, restore the non-SPS SQL databases (Master, Model, MSDB, Pubs, etc.).

3. Disconnect from the SharePoint configuration database (required as part of the disaster recovery process) from all Web server(s), Index server(s) and Search server(s).
4. Delete all the databases that are marked as 'suspect' within SQL.
5. Delete the GUID folders (located under the SPS installation path) from both the Index server(s) and Search server(s).
6. Remove SharePoint Services from the virtual server(s) from each Web server.
7. Recreate the configuration database with the same name as before and in the same location.
8. Connect to the new configuration database from the remaining Web, Index and Search server(s).
9. Configure the SharePoint topology as it was before.
Note: If separate SQL instances (different than one used by configuration database) were used for the Content and/or Component Settings databases, then configure SharePoint to point to the other instance(s).
10. From Backup Exec, select and restore all SharePoint site databases except for any Index databases.
11. From Backup Exec, select and restore the configuration database. You may need to reconnect to the configuration database from Web, Index and Search server(s) after the restore.
12. From the applicable web servers, extend the virtual servers to point (map) to the appropriate existing virtual servers (sites restored in step #10).
13. From Backup Exec, select and restore the Index databases for the applicable sites.
14. From Backup Exec, select and restore the Single Sign-on database (if applicable).

SINGLE SIGN-ON SERVICES DATABASE

If the Microsoft SharePoint Portal Farm uses Single Sign-On Services, that database as well as the Manage encryption key must be restored. The database and encryption key can be restored using the Agent for SPS Server. The Single Sign-On Services database is given the name SSO by default. The name can be configured at Single Sign-On Services setup time.

SUMMARY

The Backup Exec 10 *for Windows Servers* offers industry-leading support of Microsoft Office SharePoint Portal Server 2003. The Backup Exec Agent for SharePoint Portal Server 2003 supports SharePoint Portal Server 2003 and 2001 configurations. The agent lets you restore individual servers or an entire SharePoint Portal Server farm. In addition, individual components and database can be restored in select configurations.

VERITAS Software Corporation
Corporate Headquarters
350 Ellis Street
Mountain View, CA 94043
650-527-8000 or 866-837-4827

For additional information about VERITAS Software, its products, or the location of an office near you, please call our corporate headquarters or visit our Web site at www.veritas.com.