Microsoft® System Center

Data Protection for the Hybrid Cloud

Shreesh Dubey, Vijay Tandra Sistla, Shivam Garg, Aashish Ramdas
Mitch Tulloch, Series Editor

PUBLISHED BY
Microsoft Press
A division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Library of Congress Control Number
ISBN: 978-0-7356-9583-2

Printed and bound in the United States of America.

First Printing

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Support at mspinput@microsoft.com. Please tell us what you think of this book at http://aka.ms/tellpress.

Acquisitions Editor: Karen Szall
Developmental Editor: Karen Szall
Editorial Production: Megan Smith-Creed
Copyeditor: Megan Smith-Creed
Cover: Twist Creative • Seattle

Contents

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

http://aka.ms/tellpress

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

http://aka.ms/tellpress

Introduction

If you are responsible for architecting and designing the backup strategy for your organization, especially if you're looking for ways to incorporate cloud backup into your business continuity scenarios, this book is for you. With the increasing trends in virtualization as well as the move to the pubic cloud, IT organizations are headed toward a world where data and applications run in on-premises private clouds as well as in the public cloud. This has key implications for data protection strategy, and it is important to choose the solution that provides the same level of data protection you have afforded so far while allowing you to harness the power of the public cloud.

We will cover how the Azure Backup service has evolved into a first-class platform-as-a-service (PaaS) service in Microsoft Azure that integrates with the on-premises enterprise class backup product, System Center Data Protection Manager (DPM), to provide a seamless hybrid cloud backup solution. Current backup products treat the cloud as a storage endpoint, which we see as a limited-use case for the public cloud. The approach we describe in this book allows you to exploit the full power of the public cloud and gives you the flexibility to manage your backups in a hybrid world.

We have made a steady set of investments in DPM over the last 18 months, and, as of this writing, we have released six update rollups, including customer hot fixes as well as new features in the areas of private cloud protection, storage optimization, and workload support. The last chapter focusses on the most recently released protection for infrastructure-as-a-service (IaaS) virtual machines, which went to preview release in March 2015 and is expected to be generally available by Q3 of calendar year 2015.

This book covers improvements added in DPM 2012 R2 as well as the integration with Microsoft Azure Backup service and assumes you have working knowledge of the DPM 2012 version. To get familiar with older versions of DPM, refer to the following:

- *http://social.technet.microsoft.com/wiki/contents/articles/7485.system-center-data-protection-manager-2012.aspx*
- *http://social.technet.microsoft.com/wiki/contents/articles/11867.system-center-2012-data-protection-manager-survival-guide.aspx*
- *http://blogs.technet.com/b/dpm/*

Acknowledgments

The authors would like to thank the following individuals for their help on this book project:

- The entire hybrid cloud backup engineering team, without whom the products this book talks about wouldn't be possible
- John Loveall and the Windows Server Deduplication team for excellent collaboration to make de-duplication work for DPM
- Vijay Tewari, Jim Pinkerton, and other folks on the CPS team for making DPM the certified backup product for CPS
- Corey Sanders and Guy Bowerman from the IaaS VM team for collaborating on the IaaS VM backup feature
- Mitch Tulloch for being a great coach, for being patient about the numerous delays and schedule resets, and above all, for helping us build a book that we are really proud of and one we think will really help the community on backup

Free ebooks from Microsoft Press

From technical overviews to in-depth information on special topics, the free ebooks from Microsoft Press cover a wide range of topics. These ebooks are available in PDF, EPUB, and Mobi for Kindle formats, ready for you to download at:

http://aka.ms/mspressfree

Check back often to see what is new!

Errata, updates, & book support

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates to this book—in the form of a list of submitted errata and their related corrections—at:

http://aka.ms/Scdatapro/errata

If you discover an error that is not already listed, please submit it to us at the same page.

If you need additional support, email Microsoft Press Book Support at mspinput@microsoft.com.

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to *http://support.microsoft.com*.

We want to hear from you

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

http://aka.ms/tellpress

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

Stay in touch

Let's keep the conversation going! We're on Twitter: *http://twitter.com/MicrosoftPress*.

CHAPTER 3

Protecting Microsoft workloads

Most organizations rely on Microsoft server workloads to run their businesses. Ensuring protection of Microsoft workloads is a critical part of the business continuity strategy, and organizations need a backup tool that ensures their workload data is protected from various kinds of data-loss scenarios. With Microsoft System Center Data Protection Manager (DPM), organizations can protect workloads such as Microsoft SQL Server, Microsoft SharePoint Server, Microsoft Exchange Server, and Hyper-V virtual machines (VMs) not only to disk or tape media, but also to Microsoft Azure, Microsoft's cloud platform. This chapter describes how DPM orchestrates backup and recovery of Microsoft workloads.

Basic configuration

To start protecting Microsoft workloads with DPM, it is essential to get a high-level overview of how protection is configured. DPM discovers the set of data sources to protect, enables the selection of the data sources, and protects the data to disk, tape, or Azure. DPM also orchestrates the recovery of Microsoft workloads from the backup media to the production server or an alternate location as specified.

After DPM is installed, there are essentially three steps for configuring protection:

- Configure the replica storage to store all the backup data.
- Install DPM agents on each of the servers that need to be protected.
- Create a protection group to pick the data sources and configure the backup schedule and backup storage target for the protection group.

Adding disks to a DPM storage pool

To take advantage of built-in Windows Server de-duplication capability, it is recommended that you deploy the DPM server in a virtualized environment. In a virtualized environment, virtual disks can be added to the VM from the Hyper-V Manager. Add the disk to a storage pool using the instructions available at *https://technet.microsoft.com/en-in/library/hh758075.aspx*.

Installing DPM agents

If the servers to be protected are in the same domain as DPM, the agent install can be pushed from the DPM administrator console. If the servers to be protected are in a different domain or behind a firewall, the agent must be installed separately on each server and attached to the DPM server. Complete the following steps to install the agent from the DPM administrator console:

1. Open DPM.
2. Click the Management tab located at the bottom-left corner of the console, and then click the Agents hyperlink.
3. Click Install to launch the Protection Agent Installation Wizard. There are two options: Install Agents or Attach Agents (for servers that are behind the firewall or for computers in a workgroup or untrusted domain). Click Install Agents and select the set of target servers on which to install the agents. If the workload to be protected spans multiple servers or a cluster, the agent must be installed on each server or node of the cluster.

4. Provide inputs such as user name, password, and domain name. The user must have administrator privileges on the target computer.
5. Click Restart Method and complete the agent installation.

See also *For agent installation in different domains, follow the steps outlined in the TechNet documentation at https://technet.microsoft.com/en-us/library/bb870934.aspx.*

Configuring protection groups

With the agents installed on the servers that need to be protected, you can configure protection for the workloads. Complete the following steps to configure protection groups.

1. Open DPM.
2. Click the Protection tab, and then click New at the top-left corner.
3. Select the server to protect. DPM uses Active Directory to identify the server and the clusters in the domain that it is part of.
4. When the server is identified, select the specific set of data sources, referred to as members, to back up. You can select multiple types of data sources in a single protection group, but it is recommended that you segregate data sources based on their type and protection goals since a protection group is a means to logically group data sources that have the same protection intent.

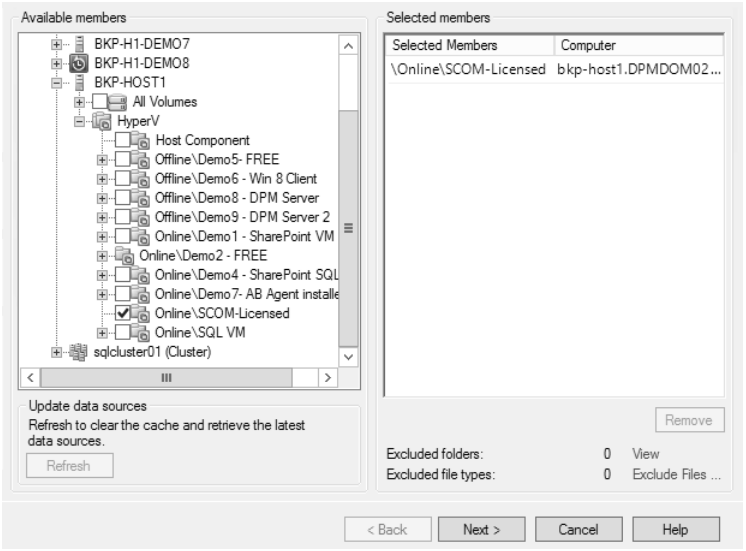

5. Select the data protection method. DPM supports short-term retention to disk as well as to tape. For short-term retention, specify the retention range as well as synchronization frequency. For long-term retention goals, select tape or Azure Backup. If a tape library is configured, the tape is shown as a valid target for long-term backup. Similarly, if Azure Backup is configured, you can select Azure as an online protection target. Subsequent chapters in this book cover how Azure Backup can be configured and used as a long-term retention target.
6. DPM carves out replica storage for each protection group. Select Co-Locate Data Source On The Same Disk if the type of data is a Hyper-V VM or client computer or SQL Server databases to gain storage efficiencies.

7. After the storage pool is selected, select the mechanism to transfer the initial copy of the data to the DPM server.
8. Select the consistency check frequency. A consistency check can run only when the replicas become inconsistent or on a daily schedule. With the latest update rollup improvements in DPM 2012 R2, the maximum duration for a consistency check job can be specified. Because the consistency check job consumes additional IOPS in the production machine, this capability is critical for administrators who want to limit the spill-over of the consistency check job past the backup window.
9. After all inputs are provided for short-term protection, a new protection group is configured and can be monitored from the Monitoring tab.
10. The protection group can be modified later to add new data sources, to modify the backup schedule or retention policy, or to add or modify long-term protection goals.

NOTE **When you select data sources, it is critical that you select the right set of members for enabling protection. For instance, when selecting clustered resources, do not pick data sources from individual nodes; instead, point to the cluster to select data sources. Similarly, when you select a SharePoint data source, it is critical to point the selection to the web front end server as opposed to the back end SQL Server instance machine.**

VSS framework

DPM leverages Volume Shadow Copy Services (VSS) and filter bitmaps to make the backup process efficient. DPM leverages file filter technology to maintain a bitmap of changes between two synchronization events. With VSS volume snapshots, the set of changed blocks as tracked by the bitmap are read and transferred from the production server to DPM replica storage. All major Microsoft workloads, such as Hyper-V, SQL Server, Exchange, and SharePoint, support VSS writers, and the Generic VSS writer enables third parties to participate in the VSS protocol. Figure 3-1 illustrates the VSS architecture.

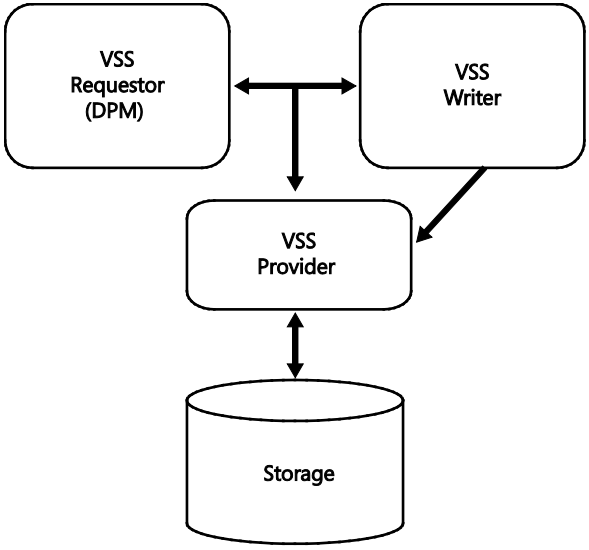

FIGURE 3-1 VSS architecture diagram

Hyper-V protection

DPM 2012 R2 UR3 Hyper-V protection features improved scalability and reliability to meet customer backup SLA. VMs can be protected to disk, to tape, or to Azure Backup for long-term retention. Hyper-V protection includes protection of Windows as well as Linux VMs in both standalone and clustered environments. Details on how Hyper-V protection is configured are described in subsequent chapters.

In a clustered environment, a protected VM can move to another Hyper-V host within the same cluster without its storage, or the VM's storage can migrate without the VM compute node, or both the compute and the storage can migrate.

With System Center Virtual Machine Manager (VMM) integration, DPM is able to continue protection of a VM during live migration without requiring user intervention as long as the DPM agent is installed on all the target Hyper-V hosts.

1. Install the VMM console on the DPM server and associate the VMM server with the DPM server from DPM PowerShell running in Administrator mode:

   ```
   Set-DPMGlobalProperty -DPMServerName <> -KnownVMMServers <>
   ```

2. Start the DPM-VMM Helper Service from the control panel.

When the protected VM's storage is migrated, a consistency check is required to make the data on the production server consistent with the data on the replica storage. This is because the file filter tracking information on the production server is not migrated along with the protected VM's VHD files, thereby causing a break in the tracking logic.

Hyper-V backup process

The Hyper-V backup process includes host-side features as well as guest-side features (see Figure 3-2).

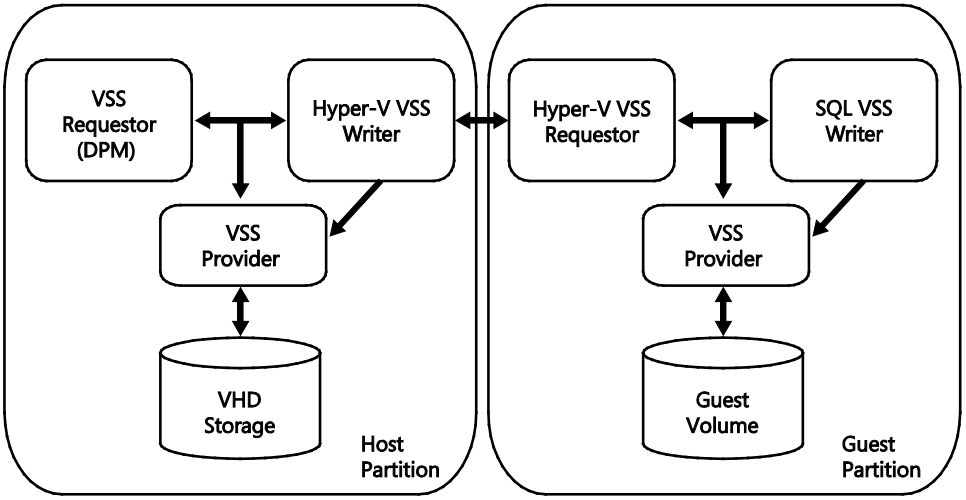

FIGURE 3-2 Hyper-V backup process

The host side includes Hyper-V VSS writer, the VSS provider, and the backup application, which acts as the VSS requestor. DPM acts as the VSS requestor, which initiates the backup operation periodically and triggers the Hyper-V VSS writer to quiesce the VM.

To enable an application-consistent snapshot of the application running within the VM, Hyper-V communicates to the guest operating system through the Hyper-V integration service. Hyper-V inside the guest acts as the VSS requestor and requests the workloads to quiesce. For instance, if SQL Server is running inside the VM, the SQL Server writer participates in the VSS protocol and flushes in-flight data buffers to the disk, and when it is done, the VSS provider takes a volume snapshot.

After the volume snapshots have been created inside the guest, the VSS provider on the host creates a shadow copy of the volume that contains the VHDs. The volume snapshot enables the application, the VM in this case, to continue to make changes to the VHDs that are attached to the VM while the backup operation takes place.

DPM maintains a bitmap file per VHD that is being tracked, and it can easily identify the blocks that have been modified since the last synchronization event. The block size is maintained as 16 KB to optimize the amount of data that is transferred and stored on the replica. Since the bitmap only indicates which 16-KB blocks are modified, the DPM backup agent reads the modified data from the storage snapshot. These changes are then transported to the DPM replica server and stored on the replica volume. To maintain versions of backup data, DPM maintains recovery points, which are essentially snapshots of the data on the replica storage pool. When a new recovery point is created on the DPM replica storage, a volume

snapshot is taken, and only new changes are written to the replica volume, while shadow copies are maintained for older recovery points. This mechanism of synchronization is called an express full backup.

Because DPM maintains a bitmap for each file that is being backed up, it requires a list of files to back up. The Hyper-V VSS writer not only participates in the backup operation, but it also enables the VSS requestor to track the list of files that need to be backed up. For instance, if a VM has two VHDs, the file specification with VSS will maintain a list of VHDs and their locations. This not only enables DPM to maintain the right tracking mechanism, but also enables DPM to capture the corresponding volume on the host machine. This enables the backup operation to detect any changes in the file specification. For instance, when a new VHD file is attached to the VM, DPM automatically makes an initial copy of the VHD as part of the next synchronization event.

During the restore process, an administrator can recover the entire VM or restore files and folders that are within the VM. For a VM, the restore can be to the same host, to the same cluster, or to a different cluster.

To ensure consistent backups are captured for Linux VMs, it is critical to install Linux Integration Services. The mechanism to quiesce the workloads within the guest operating system exists only in a Windows operating system. Hence, for Linux VMs, the Hyper-V VSS writer uses a different approach to quiesce the workloads. The Hyper-V writer leverages the file system-level operations, such as freeze and flush, to ensure that the data is file-system consistent. Therefore, from a user's perspective, the file-level consistency is always maintained by the backup process.

Client protection

Since DPM 2010, DPM has enabled protection of client computers, and with DPM 2012 R2, the latest versions of Windows 8.0 and Windows 8.1 can be protected too. With the DPM 2012 R2 latest update rollup, a backup administrator can configure protection groups to back up client data to Azure Backup. DPM enables on-the-go customers to back up their data using DPM servers when they are connected to the corporate network through wired or wireless LAN.

While the laptop or the desktop computer is disconnected from the network, the data changes to the files and folders are tracked and stored on the local hard drive of the client computer. When the client machine is connected to the corporate network, only delta changes since the most recent backup are synchronized to DPM, thereby ensuring efficient storage of incremental data changes on the DPM replica storage.

If the client computer fails to synchronize to the DPM server within the policy definition, the local DPM client user interface notifies the user to connect to the corporate network. When configuring backup of client data, an administrator can define standard folders to back up, such as My Documents, and also enable end users to add additional folders for protection. The data is automatically backed up to the DPM server at the set backup frequency.

Client data recovery

As with all other Microsoft workloads, data can be restored using the DPM console. However, the most common mechanism for restoring data is self-service recovery. After the data is backed up, users can browse and restore the recovery points using the Windows client DPM applet on the client computer. Users can search previous versions of files or browse the recovery points on a particular DPM server for all computers that they have access to.

To configure client end-user recovery, the computer on which end user recovery is desired, as well as the DPM server, must be registered to the Active Directory using DPMADSchemaExtension.exe. The following TechNet article describes the specific steps for configuring Active Directory for end-user recovery and for recovering file data: *https://technet.microsoft.com/en-us/library/jj627988.aspx*.

Exchange Server protection

With Exchange Server, customers enjoy 14 days of backup with database availability groups. However, DPM provides an excellent way to retain data for an even longer period of time. DPM enables protection to tape or to Azure for long-term backup archiving.

DPM 2012 R2 and DPM 2012 SP1 support backup of Exchange Server 2013, Exchange Server 2010, and Exchange Server 2007. With Exchange Server 2013, DPM supports backup of not only servers running Exchange but also databases configured in a database availability group (DAG). Each node of a DAG can be backed up individually with the same DPM server or with a different DPM server.

With Exchange recovery, DPM can recover a single mailbox. The mailbox database is recovered to a recovery database, and then the individual mailbox is recovered. An entire Exchange database or the entire Exchange server can be recovered if it is protected with bare metal backup.

> ***See also*** *For detailed steps to configuring an Exchange database from DPM backup, see the TechNet article at https://technet.microsoft.com/en-us/library/jj628013.aspx.*

SQL Server protection

DPM enables protection of SQL Server in various configurations. Typical SQL server configurations include standalone SQL Server, SQL Server deployed in a clustered environment with Windows Server Failover Clustering as well as SQL Server deployed as an AlwaysOn availability group. With the latest update rollup, DPM 2012 R2 also supports SQL Server 2014 along with all other major versions of SQL Server, such as SQL Server 2012 and SQL Server 2008, on all major Windows Server versions. In earlier versions of SQL Server, such as SQL Server 2008 where SQL AlwaysOn technology is not available, DPM enables protection of SQL Server in a database mirroring configuration. In a mirroring configuration, protection of the

principal database is available, although not on the mirrored database. DPM doesn't support SQL Server backup when it has database files stored on a remote SMB file share or Windows scale-out file server. DPM also doesn't support databases whose data is stored on Windows Azure blob storage.

SQL Server can be deployed in a physical machine or inside a VM. The DPM agent must be installed on the SQL Server machine. If SQL Server is configured in a clustered mode or as a SQL AlwaysOn cluster, the DPM agent needs to be installed on all the nodes that are failover targets for the SQL Server instance. If cluster members are added, DPM needs to be installed on the newly added cluster nodes as well.

In SQL Server AlwaysOn deployments, DPM honors Preferred Replica, Replica Only, Any Replica, and Primary preferences set by the SQL Server administrator (see Figure 3-3). However, for Preferred Replica, DPM always backs up only from the replica node. When Availability Group is selected for protection, all databases that are added to the availability group are automatically backed up.

> ***See also*** *For a description of how to configure protection of SQL Server AlwaysOn configuration, see https://technet.microsoft.com/en-us/library/hh780998.aspx.*

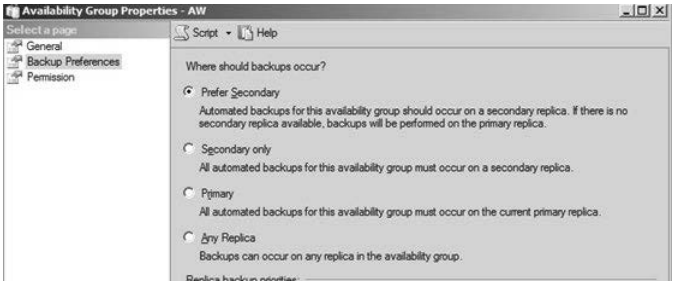

FIGURE 3-3 SQL Server AlwaysOn configuration showing backup preferences

In a SQL Server deployment, a user can enable auto protection of all databases within the SQL Server instance. This enables managing backup in a dynamic environment where databases are added or deleted from a SQL Server instance without requiring backup administrator intervention. In an auto protection mode, there is no mechanism to turn off backup of subset of databases, for example, a master or model database. Also, with auto protection, it is important to manage the size of the replica storage pool so that it doesn't run out of space.

During the restore process, an administrator can take one of the following actions:

- Recover the entire database to the original SQL Server instance
- Recover the database and rename it
- Recover to an alternate database instance
- Copy the database backup files to a restore folder in a network share

SQL Server backup process

Similar to Hyper-V VM backup, DPM leverages VSS technology to take application-consistent snapshots of SQL Server. The VSS file specification provides a list of all the .mdf, .ndf, and .ldf files that are associated for a given database that is enabled for backup. DPM maintains a bitmap file filter that tracks the blocks of SQL Server database files that are changed. For each synchronization event, SQL Server is quiesced, a volume snapshot is taken, and a stable point in time of .mdf, .ndf, and .ldf files are copied to the replica server. With DPM file filters and the change bitmap, only delta changes between the current synchronization and previous synchronization are copied to the DPM replica. With express full technology, customers can essentially do a full back up every day efficiently, both in terms of data transfer as well as storage on the replica server.

In addition, to express full backup, DPM ships a transaction log to the DPM replica storage, thereby minimizing data loss up to 15 minutes. While express full backup is efficient in terms of data transfer and storage on the replica, it is expensive on the disk IOPS since storage snapshots are maintained on the production server while the backup data is being copied. Transaction log backup, alternatively, is lightweight and enables up to a 15-minute recovery point objective (RPO).

NOTE **When you use DPM, it is critical that no other process backing or truncating transaction logs is enabled because it would interfere with DPM. After the transaction logs are backed up, they are truncated, and this could lead to a break in the transaction log chain.**

Self-service recovery of SQL Server databases

DPM can be configured to enable self-service recovery for a group of users. The first step is to configure a DPM role using the Configure Self Service Recovery option in the Protection view in the DPM console. The list of users that are allowed to perform self-service recovery are added to the role from Active Directory. The specific set of databases that are allowed for self-service recovery are added to the DPM role. The target recovery SQL Server instances are added to the DPM role.

When this option is configured, the end users can install the Self-Service Recovery Tool (SSRT) from System Center 2012 on their client computers and perform SQL Server database restore without the intervention of the backup administrator.

See also *For more information on the specific steps to configure self-service recovery and how to use the tool, see https://technet.microsoft.com/en-us/library/jj674322.aspx.*

SharePoint protection

At a high level, SharePoint consists of front-end web servers, a SharePoint configuration database, and SharePoint content databases. The main goal of protecting a SharePoint farm is to protect the content that is stored in the SQL Server content database, as well as the configuration of the SharePoint farm so that the SharePoint farm can be recovered in the event of a disaster, data loss, or corruption.

To protect a SharePoint farm, complete the following steps:

1. Install the DPM agent on the front-end servers and on each of the SQL Server instances that back the SharePoint farm. If the SQL Server instance is configured in AlwaysOn configuration, install the DPM agents on each of the servers that span the availability group.
2. Configure the front-end server for SharePoint protection using the ConfigureSharePoint cmdlet:

   ```
   ConfigureSharePoint [-EnableSharePointProtection] [-EnableSPSearchProtection] [-ResolveAllSQLAliases] [-SetTempPath <path>]
   ```

3. Use the DPM console to create a new protection group and select the member server as the front-end web server to be configured for protection.

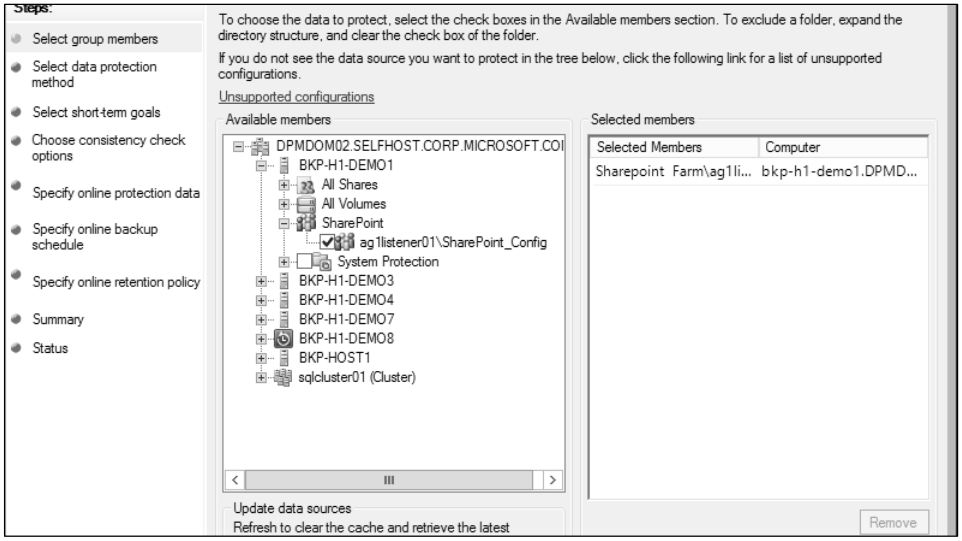

SharePoint backup process

DPM coordinates the backup across multiple servers in the farm to back up data. After backup is done, DPM queries the SharePoint object model and gets all the information about the site and items from the SharePoint server. SharePoint protection uses the SharePoint VSS writer to protect the entire farm. SharePoint VSS writer is a referential writer, and DPM uses the SharePoint VSS writer to obtain the SharePoint topology, such as the SharePoint content databases and the configuration database that are part of the SharePoint farm. With this information, the SQL Server configuration and content databases are backed up. Figure 3-4 illustrates a SharePoint 2013 configuration with backup agents.

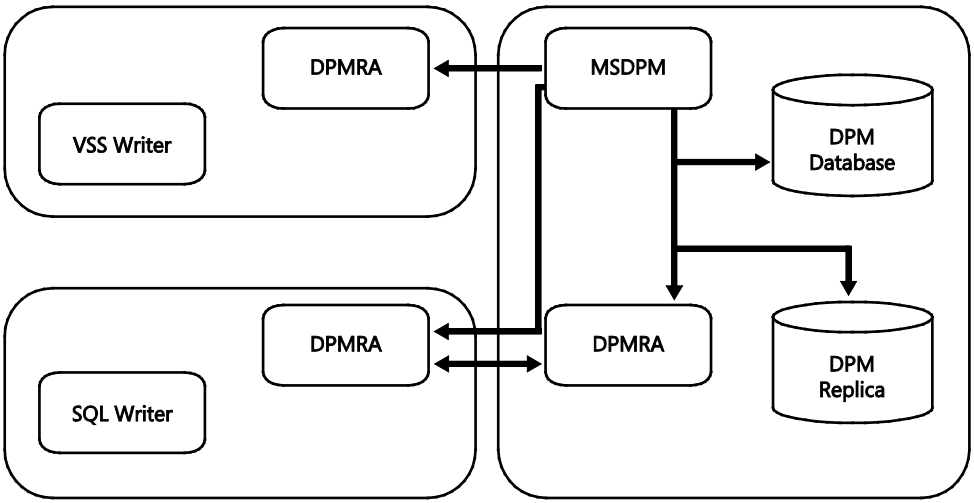

FIGURE 3-4 SharePoint 2013 configuration with backup agents

As far as SQL Server backup is concerned, it uses exactly the same workflow as backing up SQL Server databases as discussed in the previous section. The data movement for the SQL Server databases happen from the SQL Server machine to the DPM replica server directly, and the SharePoint front-end server and the VSS writer are not involved. Each content database is backed up independently since there is no referential integrity or consistency requirements across the database in a SharePoint farm.

For example, assume the SharePoint farm contains contentDB1, contentDB2, and configDB on SQL1 instance and contentDB3 and contentDB4 on SQL2 instance. SQL1 instance databases contentDB1.mdf, contentDB1.ldf, contentDB2.mdf, contentDB2.ldf, and configDB.mdf are backed up directly from SQL Server SQL1 machine, whereas contentDB3.mdf, contentDB3.ldf, contentDB4.mdf, and contentDB4.ldf are backed up directly from the SQL Server SQL2 machine. It is critical to note that although SQL backup is taken by the SQL VSS writer, transaction log backup is done as part of SharePoint SQL server backup.

SharePoint supports partial backup. Since SharePoint farms can be large, failure to back up a single database in the farm doesn't result in complete farm backup failure. Backups for one or more databases can fail, but a recovery point is created with all other databases. Similarly, when a new content database is added, its initial replica is seeded to the DPM replica server through a nightly job. To immediately add a content database to the protection group, you can add it manually.

NOTE **Filestream data that is local to SQL Server are backed up, but filestream data in a remote share are not backed up.**

TIP **If SharePoint configuration files need to be backed up, it is advised to back up the SharePoint server with System State backup.**

SharePoint catalog

In addition to the content database backup, DPM also maintains a catalog of all the items that are backed up as part of a recovery point. This information is used to identify what items are available for item-level recovery from a SharePoint recovery point. When the database backup is complete, the cataloging process starts, although it can be scheduled to run separately. The cataloging process obtains the list of all the SharePoint items that are part of the last recovery point. There is one catalog across all the recovery points in the DPM database. Failure to query the object model for the catalog doesn't result in failure of the backup operation.

Scoped consistency check

Since SharePoint farms can be large, SharePoint backup supports scoped consistency check. This enables consistency checks only on databases that are inconsistent and other databases go through normal synchronization. The recovery point is created on the DPM replica when all databases are synchronized regardless of whether a consistency check ran on some databases while normal synchronization took place on others.

SharePoint recovery

DPM enables various levels of SharePoint recovery. Since the entire SharePoint farm, including the configuration and content databases, is backed up, different parts of the farm can be restored as follows (see Figure 3-5):

- Restore entire farm
- Restore specific content databases
- Restore site collection
- Restore site
- Restore individual documents

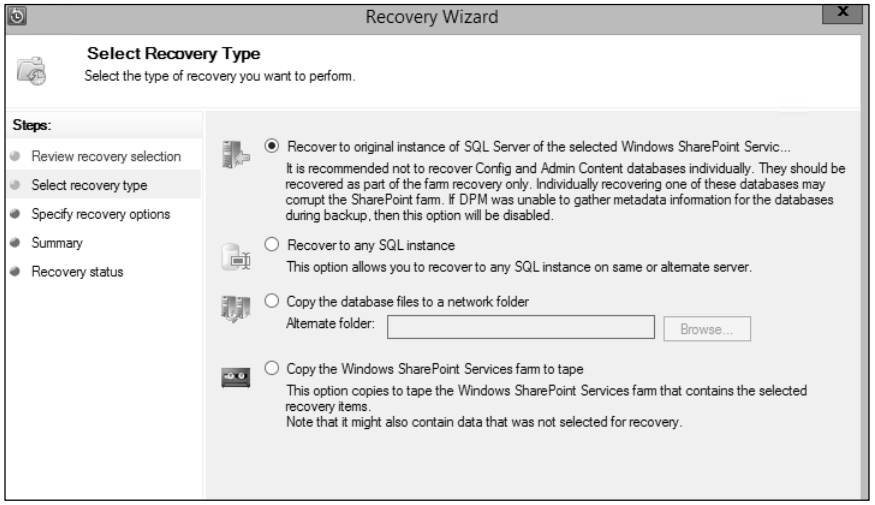

FIGURE 3-5 SharePoint 2013 Recovery Wizard

The SharePoint data can be recovered to the original site or an alternate site, or it can be copied to a network folder or to a tape. Restoring an entire site to the original site is not recommended without restoring the entire farm since it could lead to inconsistencies. The recommended practice is to recover to an alternate site or an alternate server for configuration or content database recovery. However, to restore a particular item only, it is safe to restore to the original site since the entire content is not impacted by the restore operation. If the entire site needs to be recovered from scratch, it is recommended that you recover the SharePoint server from bare metal and then recover the configuration and content databases using DPM.

To recover the farm, SharePoint recovery provides two methods. One method is to recover using a recovery farm. Use this option when the SharePoint farm has changed since the particular recovery point was created. This involves creation of an alternate farm similar to the production farm called the Recovery Farm Server. To use this option, you must specify the front-end web server name of the recovery farm as well as the SQL Server instance name and a temporary location to copy the content database to (see Figure 3-6).

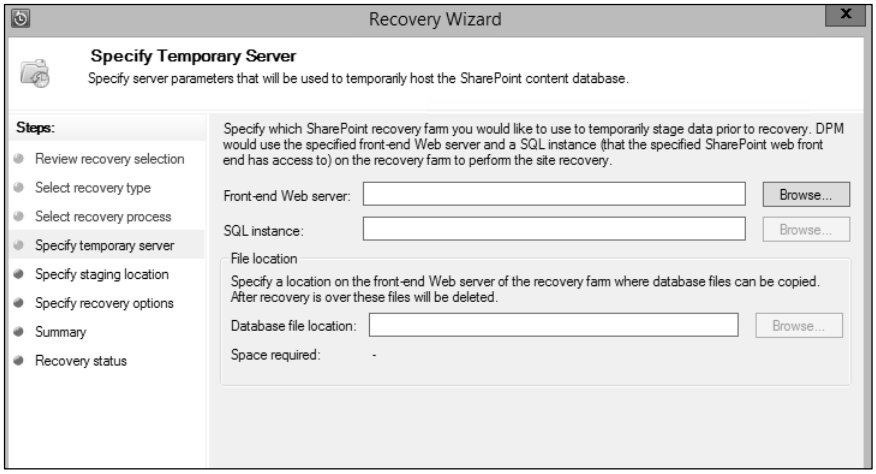

FIGURE 3-6 SharePoint 2013 recovery using a recovery farm

For item level recovery, you browse each recovery point to find the item. Information about which items are present in each recovery point and the corresponding content database mapping is included in the DPM catalog. Using the DPM catalog, the particular content database is restored first, attached to the recovery farm. Then the particular item from the content database is transferred to the original farm. Along with the contents, the URL and its security attributes are also recovered.

The other recovery method does not involve a recovery farm. This option avoids copying data to the recovery farm and then to the original farm. Instead, the SQL server contents are copied to a share and made visible to a SQL Server instance. The content database is temporarily attached to the SQL instance and then the required item is restored to the target farm (see Figure 3-7).

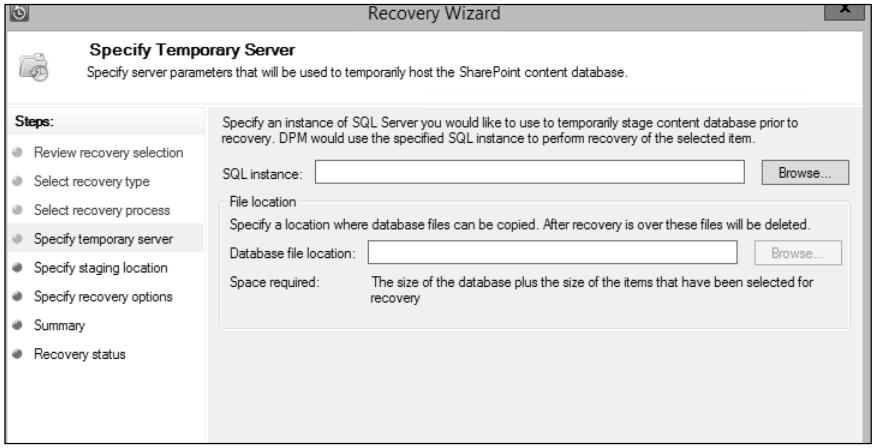

FIGURE 3-7 SharePoint 2013 recovery without using a recovery farm

About the authors

SHREESH DUBEY is Principal Group Program Manager, Enterprise Backup, for the Cloud & Enterprise Group at Microsoft Corporation. He has 27 years of experience in the IT industry (at Intel and Microsoft) with diverse technology experience in datacenter infrastructure, storage/replication, data protection, JavaScript/AJAX, developer tools, embedded systems, and EDA tools. He holds two patents in storage replication and an MSEE from the University of Texas, Austin. He currently lives in Hyderabad with his wife and two daughters, an eleventh grader and a senior at the University of California, Berkeley. He enjoys downhill skiing, all outdoor activity, cooking, and insanely spicy food.

VIJAY TANDRA SISTLA is Lead Program Manager for the Cloud & Enterprise Group at Microsoft Corporation. Vijay has 16 years of combined experience in product development, quality, and customer support in cloud, business continuity, disaster recovery, and synchronization technology areas. He holds three patents in business continuity and SQL Server technologies, and holds a Master of Mechanical Engineering degree from the University of Texas, Arlington. He is also passionate about travel and food.

SHIVAM GARG is Principal Manager – Program Management, Enterprise Backup, for the Cloud & Enterprise Group at Microsoft Corporation. Shivam has 16 years of experience in the IT industry with diverse technology experience in the data protection, replication, search relevance, local and mobile search, and supply chain domains. He holds an MBA from IIM Bangalore and a bachelor's degree from IIT Kanpur, India. He currently lives in Hyderabad, known as the City of Pearls, with his wife, Jyoti, and two kids, a fifth grader and a kindergartner. He loves water sports (snorkeling, rafting, jet skiing) and reading spiritual books.

AASHISH RAMDAS is Program Manager II, Enterprise Backup, for the Cloud & Enterprise Group at Microsoft Corporation. He has 9 years of experience in the IT industry with a balance of business and engineering functions. He started in the telecommunications domain but now holds backup/replication quite dear. He received an MBA from the Indian School of Business, Hyderabad, and a BTech in CSE from NIT Trichy, India. He is passionate about statistics and data visualization, geopolitics, pricing and COGS optimization, spicy chicken biryani, and badminton—depending on the time of the day.

About the series editor

MITCH TULLOCH is a well-known expert on Windows Server administration and cloud computing technologies. He has published hundreds of articles on a wide variety of technology sites and has written, contributed to or been series editor for over 50 books. Mitch is one of the most popular authors at Microsoft Press—the almost two dozen ebooks on Windows Server and System Center he either wrote or was Series Editor on have been downloaded more than 2.5 million times! For a complete list of free ebooks from Microsoft Press, visit the Microsoft Virtual Academy at *http://www.microsoftvirtualacademy.com/ebooks*.

Mitch has repeatedly received Microsoft's Most Valuable Professional (MVP) award for his outstanding contributions to supporting the global IT community. He is a ten-time MVP in the technology area of Windows Server Software Packaging, Deployment & Servicing. You can find his MVP Profile page at *http://mvp.microsoft.com/en-us/mvp/Mitch%20Tulloch-21182*.

Mitch is also Senior Editor of WServerNews, a weekly newsletter focused on system admin and security issues for the Windows Server platform. With almost 100,000 IT pro subscribers worldwide, WServerNews is the most popular Windows Server–focused newsletter in the world. Visit *http://www.wservernews.com* and subscribe to WServerNews today!

Mitch also runs an IT content development business based in Winnipeg, Canada, that produces white papers and other collateral for the business decision maker (BDM) and technical decision maker (TDM) audiences. His published content ranges from white papers about Microsoft cloud technologies to reviews of third-party products designed for the Windows Server platform. Before starting his own business in 1998, Mitch worked as a Microsoft Certified Trainer (MCT) for Productivity Point.

For more information about Mitch, visit his website at *http://www.mtit.com*. You can also follow Mitch on Twitter @mitchtulloch.