

EMC REPLICATION MANAGER AND MICROSOFT SHAREPOINT

A Detailed Review

Abstract

This white paper offers an in-depth look at how EMC® Replication Manager integrates with Microsoft SharePoint. The paper provides detailed information about how Replication Manager interacts to create, mount, and restore replicas of a SharePoint farm.

November 2011

Copyright © 2011 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

VMware is a registered trademark of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number h8085.1

Table of Contents

Executive summary	4
Introduction	4
EMC Replication Manager overview	4
Replication Manager Server.....	5
Replication Manager Console	5
Replication Manager agents	6
SharePoint overview	6
Captions.....	7
SharePoint application integration	7
Storage technologies	8
Storage configuration	8
Security considerations	9
Creating replicas	9
Creating a SharePoint application set.....	9
Creating a SharePoint job	10
Mounting SharePoint replicas	11
SharePoint mount options.....	12
Restoring SharePoint Replicas	12
Affected entities and SharePoint	13
Item-level recovery with Ontrack PowerControls	13
Conclusion	13
For more information	14

Executive summary

Replication Manager dramatically simplifies the management of Microsoft SharePoint through the use of disk-based replicas to improve availability and rapid recovery of mission-critical data. Replicas can be automatically mounted to alternate servers for backup and other tasks, with no impact on the production server. This paper focuses on Replication Manager's application integration with Microsoft SharePoint.

Introduction

This white paper describes the internal aspects of EMC Replication Manager and explains how it interacts with SharePoint during application set creation, replicas, mount, and restore.

The paper answers questions about “how” Replication Manager interacts with the SharePoint farm to provide value-add functionality such as backup protection. It has been written for SharePoint administrators and other technologists who want an in-depth understanding of how Replication Manager interacts with SharePoint.

EMC Replication Manager overview

This overview provides general information about Replication Manager components and their roles as a basis for more specific discussions that follow. For a more complete overview of Replication Manager, refer to the white paper *EMC Replication Manager Version 5.0 Technology - A Detailed Review* found on EMC.com. Figure 1 shows the Replication Manager architecture and the components that reside in various parts of the system.

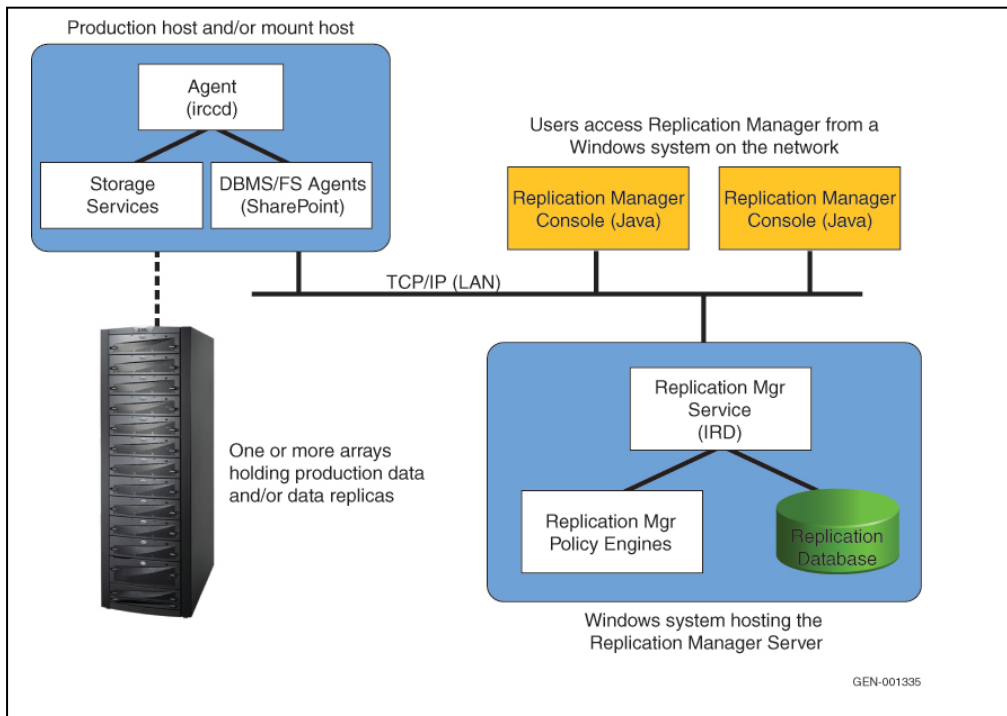


Figure 1. Replication Manager architecture

Replication Manager Server

The Replication Manager Server software is installed on a Windows system. It controls replication activities and stores data about each replica. The software has three distinct components:

- **Replication Manager Server Service** controls and coordinates replication and recovery activity for all storage corresponding to registered clients and their application sets. The service also handles all requests from the Replication Manager Console.
- **Policy Engine** links Replication Manager with the supported storage technologies. The Replication Manager Server Service uses the policy engine to analyze the storage environment and select the appropriate storage for the replica.
- **Replication Manager Repository** is an embedded database that stores data about application sets (for instance, information about the SharePoint farm), jobs (backup options, mount options, and so on), and replicas.

When installed on Microsoft Cluster Server, Replication Manager Server is a cluster-aware application and its components will fail over to a passive node in the event of a cluster failover.

Replication Manager Console

The Replication Manager Console is a portable Java application that lets you control Replication Manager from a Windows system that has a TCP/IP connection to the Replication Manager Server.

Replication Manager agents

Replication Manager agent software is installed on each host that participates in the replication process, including hosts that manage production data and hosts that are used to mount replicas for repurposing. The agent software has three distinct components:

- **Replication Manager Client Service** waits for incoming requests from the Replication Manager Server Service, then coordinates all operations on the agent.
- **Storage services component** manages the storage relationships between the client service and the storage technologies used to create the replicas.
- **Application agents** are provided for each supported file system and database application, including SharePoint. Each agent is a separate dynamic library (DLL) that is loaded by the client service as needed.

SharePoint overview

Microsoft SharePoint provides a single, integrated location where employees can efficiently collaborate with team members, find organizational resources, search for experts and corporate information, manage content and workflow, and leverage business insight to make better-informed decisions. SharePoint provides a single, integrated platform for business collaboration and has become a critical tool that must be properly protected.

A SharePoint farm is a distributed application spread across multiple Windows servers. The data resides in SQL Server databases as well as search index files on dedicated servers in the farm. Replication Manager's support for SharePoint leverages its existing support for SQL Server and for NTFS file systems. Figure 2 illustrates a typical configuration and indicates how Replication Manager would be installed in such an environment.

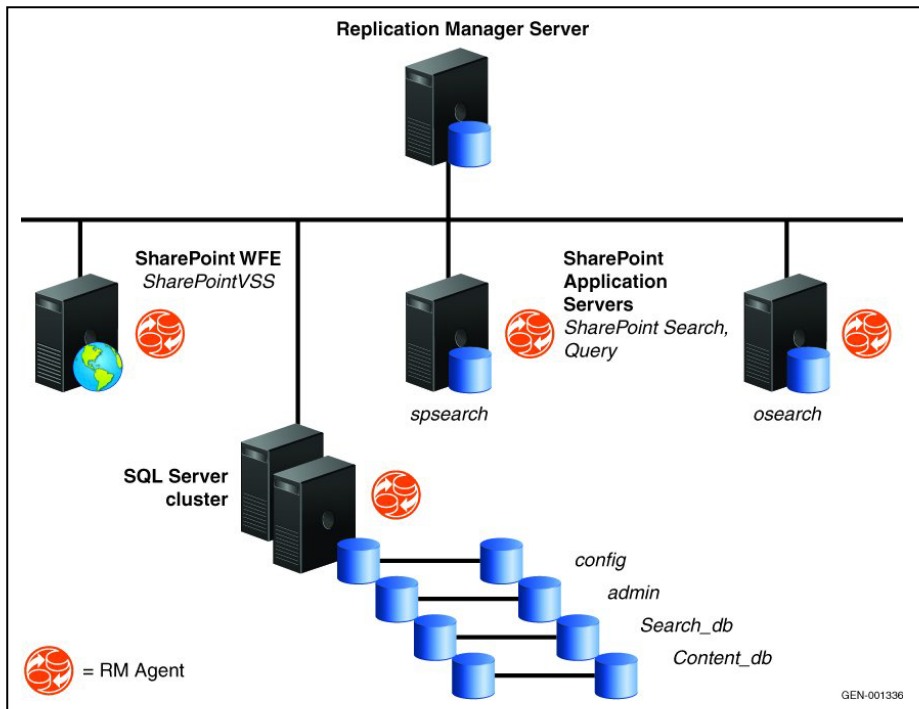


Figure 2. SharePoint 2007 configuration with Replication Manager

Typical use cases of Replication Manager’s support for SharePoint are:

- **Non-invasive backup** — Use Replication Manager to create replicas of large SharePoint farms (hundreds of gigabytes to terabytes) with minimum impact. The replicas can be used for restoring content databases or for making offline backups.
- **Disaster recovery** — Replication Manager quiesces all relevant components of the farm in an application-consistent manner so that the replica can be used for a full farm restore.

Replication Manager supports Microsoft SharePoint Server 2010 and Microsoft Office SharePoint Server 2007 SP1 and SP2. Full farm restore and content database restores are supported for SharePoint 2007. For SharePoint 2010, restore is supported at the content database level.

SharePoint application integration

Replication Manager uses the following interfaces when replicating the SharePoint farm:

- **Windows VSS framework** to obtain the layout of the farm from the SharePoint VSS Writer when creating an application set. VSS is the Volume Shadow Copy Service, a Windows service and architecture that coordinates various components to create consistent point-in-time copies of data called shadow copies.

- **SQL Server Virtual Device Interface (VDI)** to create replicas of the SQL Server databases in the SharePoint farm. Replication Manager prepares the target storage for replication by synchronizing the production with the target storage, then uses VDI to freeze I/O to the database. After the target storage is split from the production storage, VDI is used again to thaw the I/O.

Storage technologies

Replication Manager supports:


- Local replicas of SharePoint on EMC CLARiiON® storage accessed via Fibre Channel, iSCSI, Hyper-V passthrough, or VMware RDM
- Local and remote replicas of SharePoint on EMC Symmetrix® storage accessed via Fibre Channel, Hyper-V passthrough, or VMware RDM.

Storage configuration

It is important that SharePoint storage be configured so that separate LUNs are used for configuration databases, search databases, search indexes, and content databases. This is necessary as certain databases and indexes are restored at specific points during the restore procedure.

- Content databases should be on one or more LUNs that are separate from any other database or index.
- Search indexes should be on one or more LUNs that are separate from any of the SQL Server databases.
- The search databases and Shared Service Provider (SSP) database should not be on the same LUN as other databases or indexes. Note that for SharePoint 2007, a farm with multiple Shared Services Providers (SSPs) enabled is not supported.

Figure 3 shows a typical configuration.



GEN-001337

Figure 3. SharePoint storage configuration

Security considerations

Replication Manager and SharePoint require certain permissions and rights to configure application sets, run jobs, and perform restores.

Establish a common set of credentials for all SQL Server instances used by the SharePoint servers. Configure the user account to use either SQL Server authentication or Windows authentication. The account can either be a member of the local Administrators group or a non-Administrator account.

The SharePoint farm account (also known as the server farm account) must have local administrative rights on the SharePoint VSS Writer host.

Creating replicas

Replication Manager creates full replicas of an entire farm (content databases, configuration database, administration database, search databases, search index files). Partial replicas are not supported at the farm level.

Creating a SharePoint application set

An application set defines the part of a SharePoint farm that is to be replicated. When you create an application set, Replication Manager discovers the layout of the SharePoint farm as follows:

1. Connects to a SharePoint host (lrma091 in the following sample) that is running the SharePoint VSS Writer.
2. Uses the VSS framework to obtain the layout of the farm from the SharePoint VSS Writer.
3. Queries the user for the SQL Server credentials and SharePoint farm credentials.
4. Validates the credentials and displays the layout of the farm.

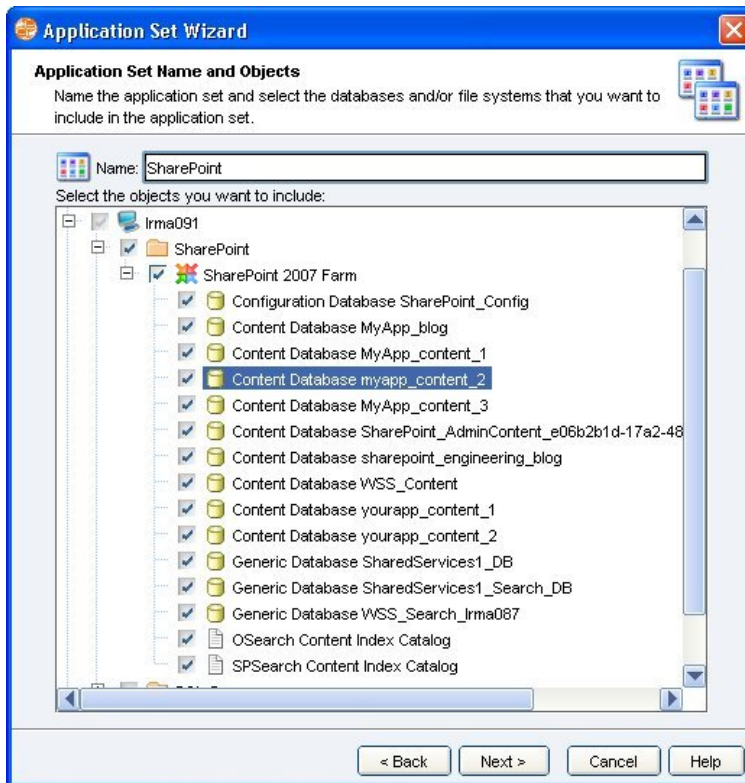


Figure 4. Selecting a SharePoint farm in the Application Set Wizard

The user sees a single application set (Figure 4), but internally Replication Manager creates one application set per host in the farm. The internal application sets are used when running a job to replicate the entire farm. They are:

- One application set containing all the SharePoint databases on the SharePoint host.
- If applicable, an application set containing all the SharePoint search indexes (Office Search, SharePoint Search).
- A separate application set to interact with the SharePoint VSS Writer. It is used for replication and restore, to inform SharePoint when farm elements have been backed up or restored.

Creating a SharePoint job

A job defines a set of actions, such as the creation and mounting of a replica, associated with an application set. It also specifies the replication technology for the replication, how long to save each replica (the retention period), the maximum number of replicas in a rotation, and whether to start user-defined scripts as part of the job.

For simplicity, Replication Manager generally presents a SharePoint job to the user as a single item. In fact a SharePoint job consists of a master job, one sub-job per host, and one VSS job. The master job is what the user sees and manages in the

Replication Manager Console. This is similar to the structure of a SharePoint application set.

SharePoint job creation prompts for the same configuration information as job creation for SQL Server 2005/2008 jobs. The Advanced Replication Settings panel of the job wizard for SharePoint is shown in Figure 5.

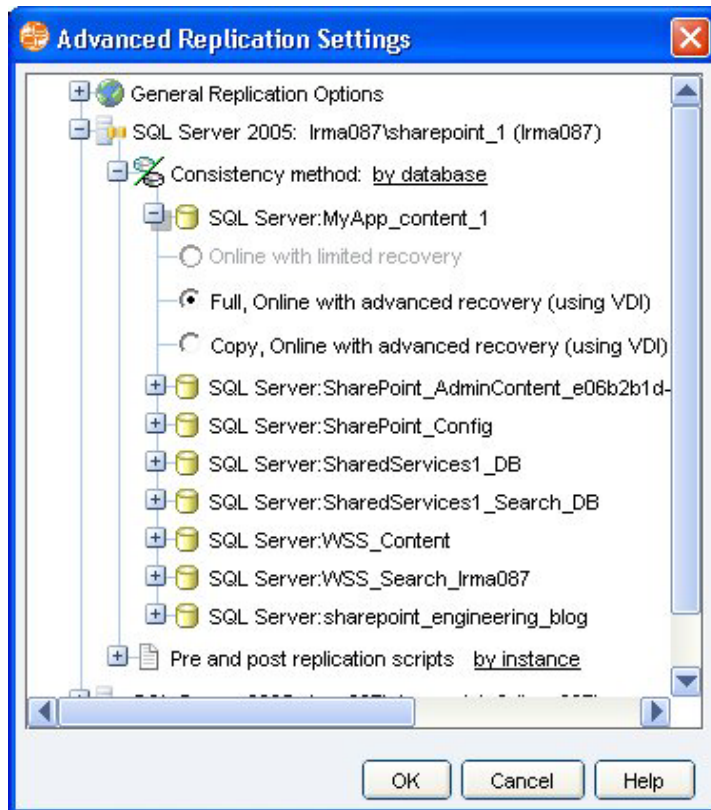


Figure 5. Advanced Replication Settings panel of the job wizard for SharePoint

Mounting SharePoint replicas

Replication Manager can perform mounts of a SharePoint replica to one or more alternate mount hosts, or in certain cases, to an alternate location on the production hosts. This section describes considerations regarding these mounts.

Replication Manager mounts SharePoint replicas at the host level.

After the target mount host has been chosen and the mount options have been selected, the Replication Manager client on the mount host:

- Performs the necessary storage technology specific operations to allow the LUNs making up the replica to be visible and ready to use on the mount host
- Mounts file systems
- Optionally performs recovery on the SQL Server databases

- Optionally runs any post-mount scripts

SharePoint mount options

Figure 6 is an example of mount options for a multi-host SharePoint setup in the Replication Manager job wizard.

The example shows the mount components of three hosts — lрма216, lрма218, and lрма219 (hereafter referred to as 216, 218, and 219).

Hosts 218 and 219 are SQL Server hosts and the standard SQL Server mount options are available. Host 216 can be identified as the search index host, because only the options for file systems are displayed.

Replication Manager can perform mounts of a SharePoint farm to one or more alternate mount hosts, or in certain cases, also mount to an alternate location on the production hosts.

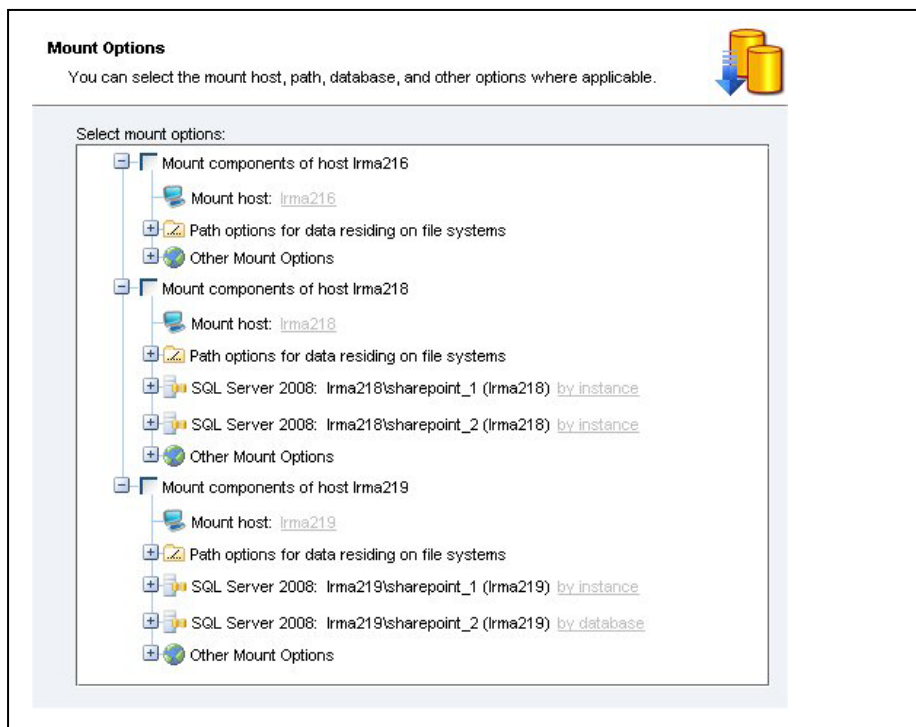


Figure 6. SharePoint mount options

Restoring SharePoint Replicas

Replication Manager can restore SharePoint replicas at the full farm level, or at the content database level. SharePoint hosts that participate in a restore are:

- the SharePoint VSS Writer host defined in the application
- all hosts that have SQL Server databases or search indexes.

As described earlier, storage for a full farm restore must be configured so that separate LUNs are used for configuration databases, search databases, search indexes, and content databases. This is necessary as databases and indexes need to be restorable at specific points during the restore procedure.

The granularity that is possible for a partial restore is determined by the number of LUNs used to back up the content databases. The most flexible configuration is one content database per LUN.

When restoring a SharePoint replica, Replication Manager will detect if multiple databases are on the same drive and will automatically select both when one is selected. You can override this selection, but Replication Manager will flag it and ask for verification of the choice before continuing.

Affected entities and SharePoint

The Replication Manager SharePoint agent analyzes the farm components and their logical layout on the storage array. It ensures discovery of the underlying SQL Server datafiles and file systems where SharePoint search index files reside. When a job runs, Replication Manager maps the production data and, based on the job option values, selects the appropriate target devices, replicates the data, and stores the mapping with the replica.

The granularity of restore operations is tightly linked to the underlying storage layout on which SharePoint farm was built. When restoring a SharePoint replica, Replication Manager will detect if multiple databases are on the same drive and will automatically select both when one is selected. You can override this selection, but Replication Manager will flag it and ask for verification of the choice before continuing.

Item-level recovery with Ontrack PowerControls

Ontrack PowerControls for SharePoint, a product of Kroll, Inc., allows administrators to do item-level recovery from SharePoint backups. No recovery farm is needed. You simply mount the Replication Manager replica, point to the replica using PowerControls, browse the content and restore the desired items.

For more information and step-by-step procedures, see the EMC Global Solutions white paper *EMC Replication Manager and Kroll Ontrack PowerControls for Granular Recovery of SharePoint Items*.

Conclusion

SharePoint is one of the fastest growing server-based applications from Microsoft. Customers are deploying SharePoint to solve their collaboration and content management challenges with SharePoint's web-based solution. EMC Replication Manager integrates with SharePoint and can provide unprecedented protection for this data.

For more information

The following resources give additional information on Replication Manager and how to use it in a SharePoint environment.

Replication Manager page on www.EMC.com

- *EMC Replication Manager Product Guide*
- *EMC Replication Manager Administrator's Guide*
- *EMC Replication Manager Release Notes*