

**Backing up Applications with
NetWorker Modules**

EMC Proven Professional Knowledge Sharing 2009



Aaron Kleinsmith
P&E Consultant, EMC Education
EMC²
kleinsmith_aaron@emc.com



Backing up Applications with NetWorker Modules

Aaron Kleinsmith
P&E Consultant, EMC Education
EMC²
kleinsmith_aaron@emc.com

Table of contents

TABLE OF CONTENTS	2
SCOPE	4
AUDIENCE	4
EXECUTIVE SUMMARY	4
APPLICATION BACKUPS	5
CONFIGURING NETWORKER FOR APPLICATION BACKUPS	6
WHICH SERVER ARE YOU BACKING UP?	8
<i>Client resource</i>	8
WHEN SHOULD THE SYSTEM BE BACKED UP?.....	8
<i>Group resource</i>	8
WHAT LEVEL OF BACKUP SHOULD BE PERFORMED?	11
<i>Schedule resource</i>	11
WHAT DO YOU WANT TO BACKUP ON THAT SYSTEM?	11
<i>Save set: and Backup command: attributes</i>	11
HOW LONG DOES THE DATA NEED TO BE RETAINED?	12
<i>Browse Policy: and Retention Policy:</i>	12
WHAT VOLUMES SHOULD BE USED FOR THE BACKUP?.....	13
<i>Pool resource</i>	13
WHAT DEVICES SHOULD BE USED FOR THE BACKUP?	14
<i>Pool resource</i>	14
WHAT OTHER DATA NEEDS TO BE BACKED UP?.....	14
<i>Save set: attribute, Directive resource</i>	14
HOW CAN WE PERFORM RECOVERY?	15
NETWORKER MODULE FOR ORACLE	16
WHICH SERVER ARE YOU BACKING UP?	17
<i>Client resource</i>	17
WHEN SHOULD THE SYSTEM BE BACKED UP?.....	17
<i>Group resource</i>	17
WHAT LEVEL OF BACKUP SHOULD BE PERFORMED?	18
<i>Schedule resource and RMAN script</i>	18
WHAT DO YOU WANT TO BACK UP ON THAT SYSTEM?	18
<i>Save set attribute in Client resource</i>	18
WHAT COMMAND SHOULD BE USED TO PERFORM THE BACKUP?	19
<i>Backup Command attribute in Client resource</i>	19
ORACLE RESTORE AND RECOVERY	20
NETWORKER MODULE FOR MICROSOFT SQL	20
WHICH SERVER ARE YOU BACKING UP?	21
<i>Client resource</i>	21
WHEN SHOULD THE SYSTEM BE BACKED UP?.....	21
GROUP RESOURCE	21
WHAT LEVEL OF BACKUP SHOULD BE PERFORMED?	21
<i>Schedule resource</i>	21
WHAT DO YOU WANT TO BACK UP ON THAT SYSTEM?	22
<i>Save set attribute in Client resource</i>	22
WHAT COMMAND SHOULD BE USED TO PERFORM THE BACKUP?	22
<i>Backup Command attribute in Client resource</i>	22
MICROSOFT SQL RESTORE	23

NETWORKER MODULE FOR MICROSOFT EXCHANGE.....	23
WHICH SERVER ARE YOU BACKING UP?	24
<i>Client resource</i>	24
WHEN SHOULD THE SYSTEM BE BACKED UP?.....	25
GROUP RESOURCE	25
WHAT LEVEL OF BACKUP SHOULD BE PERFORMED?	25
<i>Schedule resource</i>	25
WHAT DO YOU WANT TO BACK UP ON THAT SYSTEM?	26
<i>Save set attribute in Client resource</i>	26
WHAT COMMAND SHOULD BE USED TO PERFORM THE BACKUP?	27
<i>Backup Command attribute in Client resource</i>	27
MICROSOFT EXCHANGE RESTORE	27
NETWORKER MODULE FOR MICROSOFT APPLICATIONS	28
NETWORKER MODULE FOR SAP R/3	30
WHICH SERVER ARE YOU BACKING UP?	30
<i>Client resource</i>	30
WHEN SHOULD THE SYSTEM BE BACKED UP?.....	31
GROUP RESOURCE	31
WHAT LEVEL OF BACKUP SHOULD BE PERFORMED?	33
<i>Schedule resource or configuration file</i>	33
WHAT DO YOU WANT TO BACK UP ON THAT SYSTEM?	33
<i>Save set attribute in Client resource</i>	33
WHAT COMMAND SHOULD BE USED TO PERFORM THE BACKUP?	34
<i>Backup Command attribute in Client resource</i>	34
SAP R/3 WITH ORACLE RESTORE AND RECOVERY	35
CONCLUSION	36
REFERENCES	37
BIOGRAPHY	ERROR! BOOKMARK NOT DEFINED.

Disclaimer: The views, processes or methodologies published in this compilation are those of the authors. They do not necessarily reflect EMC Corporation's views, processes, or methodologies

Scope

This article provides NetWorker® Administrators with guidelines on how to configure and implement application and database backups. It does not replace the Installation Guide, Administration Guide or Release Notes that document each NetWorker Module. This article highlights the similarities and differences between some of the NetWorker Modules, and guides an experienced NetWorker Administrator to configure NetWorker to backup applications and databases in their environment.

This article will not address Snap-shot back up and recovery, but rather the more traditional online (hot) and offline (cold) backup methods. Also, we will not address how to backup and protect applications or databases that are protected with clustering software.

Audience

This article benefits EMC customers who are NetWorker backup administrators, EMC system integrators, or EMC employees interested in backing up applications such as Oracle, Microsoft SQL, Microsoft Exchange and SAP R/3 with Oracle. These are some of the more commonly deployed NetWorker Modules.

Executive Summary

There is information available relating to individual NetWorker Modules, but no single document offers get a complete picture of how to back up multiple types of databases likely to be found in an environment. This article will help to fill a need by addressing backup and recovery of some popular application types with NetWorker.

We will discuss traditional back up methodologies such as online (hot) backups, offline (cold) backups, and transaction log backups. This article describes the backup procedure to capture consistent copies of an online application from the source disks using the application servers' resources that manage the primary copy of the data.

We will overview the Modules with enough technical information to help you configure a backup. A storage administrator responsible to manage NetWorker could set up, start, monitor, and verify backups properly for critical application and database servers. You can use the NetWorker Client Configuration wizard to configure initial backups. However, this does not prepare you to properly respond to changing requests or questions from database administrators who manage the application and database servers. It doesn't help you to accurately report backup success and failures to managers, or to troubleshoot problems. These are important reasons to understand exactly how NetWorker Modules are configured and how they operate.

Application Backups

Current management strategies and technologies for applications have evolved to a point where it is commonly accepted that an application (application or database) will be available all the time. There are exceptions to the seven days a week, 24 hours a day (7x24) requirement, but 15-20 years ago it was common to shut down every application server when it required a backup. Some related problems included not having enough technologically advanced hardware in place, impact on server resources were unbearable during the back up, cumulative days of log backups unreasonable for recovery, or simply that the application did not need to be available 7x24. It was generally acceptable to have a window of time when the application would be unavailable.

Today, this notion is ludicrous; most applications are expected to be available 7x24. There are many technologies available to ensure that an application is always available such as clustering and replication with options to include offsite protection. However, it is still important to perform data backups to ensure that a restore can be done in the event that one of the availability solutions fails, or an application does not warrant a more costly solution than backup.

Occasionally, an application can be shut down on a daily or weekly basis making it appropriate to perform an offline backup. This kind of backup is considered a 'cold' backup.

Some application backup utilities can request an offline backup through their native backup utility. NetWorker commands such as `savenpc` (save pre-n-post command) can call custom scripts to shut down an application for backup automatically, and then restart them at the end of the backup. However, the application will be unavailable during the entire backup period.

Performing an online (hot) backup involves backing up the application while it is running and available to users. Online backups prevent us from shutting the application down during backup. NetWorker has many Modules that can perform supported online backups of different application types through supported Application Programming Interface (API) or backup utilities native to the different applications. Performing backups through NetWorker modules allows the NetWorker software to control the backup media, store save set and volume information, and perform all scheduling tasks.

Configuring NetWorker for Application Backups

This article focuses on performing online (hot) backups through the NetWorker Modules that integrate with the application vendor's API or backup utility. Use these strategies with local area network (LAN) backups or LAN-free storage area network (SAN) backups. In other words, the application server in the following scenarios could be considered a NetWorker client or a NetWorker Storage Node.

We will cover common configuration steps with the NetWorker Modules in the first part of this article. The second half covers the more detailed information and settings for each of these modules:

- NetWorker Module for Oracle
- NetWorker Module for Microsoft SQL Server
- NetWorker Module for Microsoft Exchange Server
- NetWorker Module for Microsoft Applications
- NetWorker Module for SAP R/3 with Oracle

Some of the information and guidelines are relevant not only to the five modules covered in detail, but to other NetWorker Modules that protect applications such as Lotus Domino, Informix Dynamic Server, DB2 Universal Database, Sybase Adaptive Enterprise Server, MEDITECH and Documentum.

It is important to understand business requirements before you configure a backup for your application. You may find clear and concise backup policies defined in your organization, but they may not match to the NetWorker configuration tasks, or requirements may have changed since the policies were initially written.

The following table shows some common questions to help you configure a NetWorker Module backup:

Business requirements	NetWorker Resource to be configured
Which server are you backing up?	Client resource
When should the system be backed up?	Group resource
What level of backup should be performed?	Schedule resource or configuration file
What do you want to back up on that system?	Save set attribute in Client resource
What command should you use to perform the backup?	Backup Command attribute in Client resource
How long does the data need to be retained?	Browse, Retention attributes in Client resource
What volumes should be used for the backup?	Pool resource
What devices should be used for the backup?	Pool resource
What other data needs to be backed up?	Save set attribute, Directive resource

NetWorker and the NetWorker Modules are very flexible to help backup the different types of server environments and meet the different backup requirements that come with managing these environments. With flexibility, there are usually multiple ways to fill requirements.

In the following scenarios, I'll present the easiest or most common method to configure an option to back up the applications. Please consult the appropriate Administration Guide for the NetWorker Module in question to see different methods to perform backups and recoveries.

Which Server are you Backing Up?

Client resource

When configuring an application server for backup with NetWorker, it will need multiple Client resources in the NetWorker configuration interface. The NetWorker client is the application server that needs to be backed up. Each newly created client needs identical client names to match the hostname of the application server being backed up.

We may have multiple client instances to accommodate for a client backing up file system data, a client to perform full database backups, and one or more other clients to perform transaction log or incremental backups at shorter intervals than the full database backups. Each client resource will have different properties such as backup times, different save set properties, and different NetWorker Module commands or scripts that will need to run.

We can differentiate between the different client instances by putting in different 'Comments' to describe what each client is configured to backup. This 'Comments' attribute is displayed in the Client resource listing in the NetWorker Management Console (NMC) and makes it easy to differentiate one client resource from another.

When Should the System be Backed Up?

Group resource

You can start a regular file system backup of a NetWorker client from the client through the NetWorker User GUI on a Windows client, or through the save command on other operating system (OS) types such as UNIX or Linux. This type of backup is a client-initiated or manual backup.

Most NetWorker Modules have the ability to perform manual backups. Execute a NetWorker Module or native backup utility command instead of using the save command for a backup. You can perform a manual backup on-demand from the client, or integrated with scripts and third-party schedulers to schedule database backups.

Initiate the save command through a configured Group resource on the NetWorker server for a regular scheduled backup. This type of NetWorker backup is a server-initiated or scheduled backup. In this case, NetWorker is configured to automatically begin backups whenever the Groups are scheduled to start.

We can use the same framework during a database backup with the NetWorker Modules. The NetWorker Module command is executed on the application server during the scheduled backup instead of the normal save command.

Test a file system (OS) scheduled backup first to make sure it can complete successfully before configuring a scheduled backup for your NetWorker Module configuration. There is a very strong possibility that an application scheduled backup will fail if a scheduled file system backup of a client cannot be performed.

Full database backups are typically performed at least once a day early in the backup window. There may be degraded performance during an online backup since the running database and the backup processes are sharing server resources. Although a database may offer 7x24 availability, it may have high peak usage only at certain times. Choose backup start times so the backup can be performed during low-usage times.

You can configure a new Group resource with the Start time: specified in 24 hour time and the Autostart attribute enabled for performing a full backup once a day.

You should perform transaction log (incremental) backups more frequently than full application backups. Transaction logs recover a database or application to a specific point in time. Different applications have different ways of handling incremental backups, but it is always important to back up these changes frequently. This gives you the option to recover an application to a different point-in-time than the last full backup.

One strategy is to create multiple Group resources to start incremental backups at different times of the day to perform frequent transaction log backups for any of the applications.

Create multiple Group resources, each with a different start time. Make the client responsible for calling the incremental backups a member of each group. For example, to perform a backup every 60 minutes during regular business hours, create eight different groups with each with a specific start time.

This table illustrates this:

Group Name:	Start time:	Autostart:	Interval:	Group members
LogGroup9am	9:00	Enabled	24:00	db-server.emc.com
LogGroup10am	10:00	Enabled	24:00	db-server.emc.com
LogGroup11am	11:00	Enabled	24:00	db-server.emc.com
LogGroup12pm	12:00	Enabled	24:00	db-server.emc.com
LogGroup1pm	13:00	Enabled	24:00	db-server.emc.com
LogGroup2pm	14:00	Enabled	24:00	db-server.emc.com
LogGroup3pm	15:00	Enabled	24:00	db-server.emc.com
LogGroup4pm	16:00	Enabled	24:00	db-server.emc.com

Another approach is to use a non-NetWorker scheduler (cron on UNIX, Windows Task Scheduler, EMC Control Center[®], etc.) to call the backups. Other scheduling software products can specify when a backup should start at specific time and restart at an interval, but also offer the option to include a window of time when the re-occurring backup should not run. For example, start at 9:00am and backup every 60 minutes except between the hours of 5:00pm and 9:00am the following day – when the full database backup is occurring. These scheduling programs can call the savegrp command to run on the NetWorker server for a single LogGroup at different times of day.

Some scheduling products call scripts in a hierarchical order – one after the other. For example, don't start script2 until script1 is complete. This prevents an incremental backup from starting until a full backup has completed, or prevents the file system (OS) backups from overlapping with the application backups.

If you decide to perform manual data backups, the client file indexes (CFI) of these clients will not be backed up automatically. Make sure that the CFIs are still backed up automatically through a Group resource if you choose to perform manual backups on a regular basis without scheduled backups.

What Level of Backup should be Performed?

Schedule resource

Each NetWorker Module handles NetWorker level backups differently. Most support one or more level settings in a Schedule resource when used during a scheduled backup.

All of the NetWorker Modules support some kind of incremental or transaction log backup whether set in the Schedule resource, or in a configuration or script file. Consider the amount of time it will take to apply the logs during a recovery when deciding on a schedule for incremental or log backups. A full weekly backup may be sufficient in a smaller environment with fewer transactions. You may need to perform full daily backups for a larger application server. Consider restore/recovery Service Level Agreements (SLA), and test recoveries before deciding on a scheduling strategy.

What Do you Want to Backup on that System?

Save set: and Backup command: attributes

For a file system backup, a save set will be a copy of partition, file system, directory or list of files to be backed up as specified in the Client resource "Save set:" attribute. When performing backups with the NetWorker Modules, the Save set attribute is used differently than for a file system type backup. Each Module has unique keywords and values used to indicate what type of backup is taking place.

Each NetWorker client resource has "Backup command attribute" used to perform a regular save for a file system backup, savepnpc, a script or a NetWorker Module command.

When you leave the Backup command attribute blank and a scheduled backup starts, the save command on the client is used to backup the file systems and directories listed in the "Save set:" attribute. When performing scheduled backups of the NetWorker Module clients, the Backup command attribute will contain a command or script to initiate the application backup on that client. Some of the commands require you to include a special switch to point to a configuration file.

The NetWorker Module commands, native backup utilities, and APIs used during the backup request a consistent copy of the application data. This data includes transactional data in memory and active data on disk. Backup the application data in quiesced mode (or backup mode, i.e. when the data is quiet and not active) so that critical components of the application are not changing during the backup. The NetWorker Module coordinates these backup activities.

The NetWorker Modules that require configuration files or scripts usually include a template file that you can use as a base for creating custom configuration files and scripts with your own backup environment needs. Always copy the template files to a new name and then modify the new copy. Make sure the new filename begins with either nsr or save when you are modifying a script (for example, .sh or .bat files). NetWorker will not execute any commands from the Backup command attribute that do not follow this naming convention.

How Long Does the Data Need to be Retained?

Browse Policy and Retention Policy:

A NetWorker client has two policies to help determine how long to maintain data in the indexes; the browse policy and retention policy.

The browse policy allows you to configure how long to keep backup information in the client file index. This will keep the information available for a simple, browse-able recovery. The browse policy helps the NetWorker administrator control the size of the client file indexes. Typically, client file indexes (CFI) for NetWorker Module clients will be smaller than comparable NetWorker file system clients. For example, a NetWorker Module client with a 2 TB Microsoft Exchange server will have a much smaller CFI than a comparable 2 TB file server with millions of small files.

The retention policy allows you to configure how long to protect information from being overwritten or removed from your backup volumes. Volumes that have expired save sets are considered eligible for re-use and may be overwritten with new backup data.

Typically, the application data and the file system data from the same client will have different retention requirements. This is an important consideration when deciding on what NetWorker pools to create for your backups; we will discuss in the next section.

What Volumes Should be Used for the Backup?

Pool resource

The NetWorker server uses the pool resource to sort backup data to different groups of backup volumes. A media pool is a collection of volumes (tapes, virtual tapes, and disk) that contain backup data.

Using different pools for application backups is not required, but often needed for better recycling efficiency and ease of management.

Application data can be subjected to different laws and policies than the file system data, and may be subjected to legal discovery requirements that the file system data is not. Therefore, the retention policy for the file system data is typically shorter and the application data longer.

It is not possible to recycle parts of a tape if tape media or virtual tape media is being used for backups. When everything on a particular tape expires, the whole tape is eligible to be reused and recycled. If there is a small amount of data on a tape that will not expire for a long time, the whole tape cannot be re-used until that data expires. For example, if there is a tape that is full and is 90% file system data that expires in 1 year and 10% database data that will not expire for 7 years – the 90% of the tape that will expire in one year will still not be used again (recycled) until after the database data has expired after 7 years. This is inefficient use of tape media.

An alternative is to have two separate NetWorker pools, one pool for the 1 year file system data and one pool for the 7 year application data. Separate media pools for application data also prevents database backups inter-mixed (multiplexed) with other client or file system backups. This provides quicker recovery performance with physical tape recoveries.

We can easily identify which tape volumes will need to come back on site when data has expired by sorting the data onto volumes with matching retention criteria. It will be easy to identify when the tape should return if all of the data on an off-site tape expires around the same time.

What Devices Should be Used for the Backup?

Pool resource

NetWorker pools can dedicate a particular pool to use specific backup devices. Backing up log files from a database can be a critical function to free up disk space for new transaction logs. Usually, removing a transaction log file from the disk is not acceptable until we have a backup copy. Having dedicated backup devices for log backups ensures that they are always available for backup and restore operations of log files. This way, older log files can be removed from the disk.

Virtual tape devices and disk arrays are ideal targets for transaction log backups since the random-access nature of disk storage provides quick restores. The device attribute of the Pool resource allows a pool to be directed to specific device types, such as virtual tape or disk. Creating a pool specifically for transaction log backups on these devices allows data to be directed away from the physical tape devices.

Pools for clone copies of backup data can then be written to physical tape devices creating off-site tape copies to meet any disaster recovery requirements.

What Other Data Needs to be Backed Up?

Save set: attribute, Directive resource

Some of the applications that the NetWorker Modules support store their data outside of the application. These files might be stored as regular file system data and may not be protected by the application backup. You can usually protect this data by running a typical NetWorker backup using the save command.

You can use an “All” save set for the file system (OS) backup, but the backup may try to capture file system data including active application files that reside on this NetWorker client. In the case of a Windows NetWorker client, errors will be generated when trying to access the files since they are already in use by the application. In the case of UNIX or Linux, we may be able to copy the data files, but they are not useful if they were changing at the time of backup since new information would be missing. Furthermore, it could be confusing during recovery to have them available for selection in the NetWorker recovery interface.

You can configure a NetWorker directive resource to exclude the application files during the file system backup. That way you can use an “All” save set to capture all local file systems, including any newly added file systems that may be accidentally excluded if we were listing them in the save set attribute.

How Can we Perform Recovery?

With most of the applications that the NetWorker Modules support, there is a distinct terminology difference between performing a recovery and performing a restore. A restore is what NetWorker typically refers to as a recovery or restore – simply bringing data back from the back up media. However, with the application terminology, recovery has a specific meaning. Recovery, when referring to a database, is generally defined as the application of restored transaction logs or incremental backups to the database files to bring it forward to a specific point of time.

Restore is the process that the NetWorker Module will perform to bring data back from backup media. Recovery is a process invoked by the NetWorker Module or native backup utility after data is restored back from media.

With the NetWorker Modules, database restores are performed through a NetWorker Module supplied interface or the database vendor recovery interface. The recovery is coordinated between the database API or native utility and the NetWorker server nsrd process. When the NetWorker server communicates to the NetWorker client to perform the recovery, the NetWorker Module-specific recovery command is called instead of the usual recover command (used for a file system recovery).

You must update the Remote Access: attribute in the Client resource when planning to restore data to a NetWorker client that the data was not originally backed up from. This gives privileges to the receiving client to access the original backup client's CFI.

The following sections of the article detail with specific NetWorker Modules. Each section is organized to focus on the first five questions in the Business requirements table. The NetWorker settings related to these five questions are generally different for each NetWorker Module.

Business requirements	NetWorker Resource to be configured
Which server are you backing up?	Client resource
When should the system be backed up?	Group resource
What level of backup should be performed?	Schedule resource or configuration file
What do you want to back up on that system?	Save set attribute in Client resource
What command should be used to perform the backup?	Backup Command attribute in Client resource

NetWorker Module for Oracle

You can use the NetWorker Module for Oracle (NMO) to link your NetWorker backup and recovery environment with Oracle's Recovery Manager (RMAN) to protect Oracle databases.

RMAN is the native backup and recovery utility for Oracle and is responsible for capturing consistent copies of the database for backup including datafiles, archive log files and control files. RMAN performs recovery and restore operations with the help of the Oracle RMAN recovery catalog. RMAN maintains the recovery catalog, but you should configure this data as part of your NetWorker backup outside of any steps included here.

NMO requests backups of RMAN through Oracle System Backup to Tape (SBT). Linking the NMO library to Oracle during installation means that whenever RMAN sends data to SBT it will be intercepted by NMO and passed on to NetWorker to be written to a NetWorker backup device.

Which Server are you Backing Up?

Client resource

For NMO, you must create multiple Client resources using the Oracle server hostname in order to perform different types of scheduled backups. These include full database backups, archived redo log backups, block-level incremental or differential backups and file system (OS) backup data.

When Should the System be Backed Up?

Group resource

You can start an NMO backup as a manual or scheduled backup. With a manual backup, the RMAN command is called directly on the Oracle database server from the command line or a script. Alternatively, you can perform a manual backup through the Oracle Enterprise Manager (OEM) Backup Management Tools that provide a GUI interface to RMAN.

You can perform a scheduled backup using a Group resource to configure a time that you would like the database to back up. Each type of backup requires a different RMAN script that is called through a client resource. You should perform archived redo log backups more frequently than full database backups, so you may need to create multiple groups for the different start times.

Use archived redo logs to recover a database to a specific point in time. Without archived redo logs, a database could only be recovered to the last time a full or incremental backup was performed.

What Level of Backup should be Performed?

Schedule resource and RMAN script

Only full and skip levels are supported when you configure a Schedule resource for your NMO backup. Configure the type of backup in the RMAN script. A skip backup set in the Schedule resource will not perform a backup. Full and any other NetWorker level (incr, 1,2,3,4,5,6,7,8,9) will execute the backup scripts for that day. For example, there is no real difference between a level 'Full' and level 'incr' in the NetWorker schedule, both simply tell NetWorker to execute the RMAN backup.

You can configure the RMAN script to perform any of the supported Oracle RMAN backup types including Backup full, Block level incremental, Archived redo log backups, Control file autobackup, Automatic channel allocation, Backup copies (during manual backups only), Retention policies, Backup and restore optimization, Backup of backup sets, Restartable and offline backups.

You can perform Oracle full, incremental and differential backups using appropriate RMAN settings. A level 0 backup set in the RMAN script will perform a full backup; level 1 backup will perform an incremental backup. Performing archived redo log backups is another way to capture changes on a database.

What Do you Want to Back Up on That System?

Save set attribute in Client resource

For an NMO client, the Save set attribute needs to include the full path and name of the RMAN script that will be called, preceded by the keyword RMAN: .

Two samples of the save set setting are included here:

```
RMAN:/nsr/rman_scripts/full_backup
```

```
RMAN:/nsr/rman_scripts/arch_log_backup
```

The RMAN script name is placed in the Save set attribute of the Client resource so it can be passed to the savefs command on the NMO client during a scheduled backup.

What Command Should be Used to Perform the Backup?

Backup Command attribute in Client resource

Use the Backup command attribute to specify the name of the custom nsrnmo script (UNIX, Linux) or nsrnmo.bat (Windows) file that has been configured with NMO backup settings. The nsrnmo script is used to set environment variables before the NMO nsrnmostart program is executed to request the RMAN backup. Update the nsrnmo file to include specific variables for backing up your database. The file will be named nsrnmo.bat on Windows, and nsrnmo on UNIX and Linux. The file can be found in the same location where the NetWorker executable and binary files are installed on each NMO client.

Some of the required variables that you need to modify in the nsrnmo script are included in the following table:

Environment variable	Description for value
ORACLE_HOME	Mandatory, home directory of Oracle installation
PATH	Mandatory, directory locations of nsrnmostart and save commands
LD_LIBRARY_PATH	Mandatory, Tru64 Unix only, directory location of Oracle libraries
ORACLE_SID	Required in some cases (see Admin Guide for details) ORACLE_SID used to set the SID of database to be backed up.
TNS_ADMIN	Mandatory, if Oracle Net files are not in \$ORACLE_HOME/network/admin

NMO has an additional feature, save set bundling, that you can enable to automatically create a save set bundle for each schedule backup cycle (full backup of the database object and all subsequent incremental backups) of an Oracle database object. This feature groups the dependent save sets from the same backup cycle into one save set bundle, and automatically names and records it in the NetWorker media database.

The save set bundling feature is disabled by default, but you can enable it with the `nsrnmoadmin` command.

```
nsrnmoadmin -r add NSR_BUNDLING enabled
```

Please reference the NetWorker Module for Oracle Multiplatform Version Administration Guide for more details on save set bundling.

Oracle Restore and Recovery

Start an Oracle RMAN restore by issuing an appropriate RMAN command from the RMAN command-line interface, using the RMAN command with `restore RMAN` script, or using the OEM Backup Management Tools.

We can perform restores of Oracle data files, specific table spaces, control files and archived redo log files With the NetWorker Module for Oracle.

Please reference the NetWorker Module for Oracle Multiplatform Version Administration Guide for more specific details on restore RMAN scripts and the process of restoring the Oracle database with NMO and RMAN.

NetWorker Module for Microsoft SQL

You can use the NetWorker Module for SQL (NMSQL) to link your NetWorker backup and recovery environment with Microsoft's SQL Server to protect critical application data.

The NetWorker Module for Microsoft SQL provides additional programs on the database server to perform online and transaction log backups of the SQL databases. NetWorker User for SQL Server GUI is installed with NMSQL and is used for manual backup and restores. NMSQL can perform backups of the online database including files, file groups, file streams, SQL database and transaction logs.

Which Server are you Backing Up?

Client resource

Create multiple Client resources using the SQL Server hostname to perform different types of scheduled backups (full database backups, transaction log backups and file system backup data) when performing NMSQL backups. You can also create optional clients to backup specific SQL databases at different times.

When Should the System be Backed Up?

Group resource

You can start an NMSQL backup as a manual or scheduled backup. With a manual backup, you can choose a collection of files, filegroups, databases or transaction logs from the NetWorker User for SQL Server GUI.

You can perform scheduled backups using Group resources to configure the time that you would like the database to back up. Perform transaction log backups more frequently than full database backups. You can create Multiple Group resources with different times to start transaction log backups at different times of the day. This provides a number of point-in-time options to recover the database in-between full backups.

What Level of Backup Should be Performed?

Schedule resource

There are four supported levels when configuring a Schedule for NMSQL backups: full backup, incremental backup, differential backup and skip a backup:

NetWorker level	What is backed up?
Full	Full database, then transaction logs. Logs truncated after successful back up.
1,2,3,4,5,6,7,8,9	Transaction logs backed up, not truncated. Left in place for a follow-up Full, Differential or Incr backup.
Incr	Transaction logs backed up. Logs truncated after successful backup so they are not there during a following incr backup.
Skip	Nothing is backed up– don't perform the backup

NMSQL also supports creating filegroup differential and file differential backups. A filegroup differential reduces the amount of media required for the backup and can be substituted for any log backups performed between other full and differential backups. The availability of this feature depends on the type of data that has been chosen for backup. The SQL Server must be preconfigured to enable incremental backups.

There are some variations on the supported NetWorker levels. They are documented in the NetWorker Module for Microsoft SQL Server Administration Guide if you are planning to take advantage of the filestream and filegroup backup options.

The filestream feature of SQL 2008 must be enabled on the SQL Server during a backup or restore operation.

What Do you Want to Back Up on That System?

Save set attribute in Client resource

The NetWorker client used for NMSQL backups needs a save set entry to indicate SQL is being backed up. Specify MSSQL: as the save set name if the intention is to back up all databases on the SQL Server. Alternatively, you can specify MSSQL:database1 or MSSQL:pubs to backup up a specific database within the SQL Server.

What Command Should be Used to Perform the Backup?

Backup Command attribute in Client resource

Use the Backup command attribute in the client resource to call the custom NMSQL program nsrsqlsv. This will be used to interface with the SQL Backup API and perform backups of SQL Server objects such as files, filegroups and databases. When you perform a scheduled backup of the SQL Server, place the command in the Backup command attribute of the SQL Server Client resource that is to be backed up.

Enter the command as “nsrsqlsv” in the attribute, without the quotes. If backing up SQL Server on a virtual server, you will also want to specify the “-a virtual_server_name” option used to tell NetWorker that the client is a virtual server.

Use backup striping when backing up a SQL Server to enable the use of multiple backup devices. Stripes of data can be extracted from the SQL database in parallel to enable a quicker backup by using multiple backup devices. You can enable backup striping during a manual or scheduled group backup.

Append the “-Sn” option to the snrsqlsv command to enable striping during a scheduled backup. The upper-case “S” informs NMSQL to perform the backup striping. Replace the lower-case “n” with a number to specify the appropriate number of backup stripes.

Microsoft SQL Restore

With the NMSQL, you can perform four types of restore including copy, verify only, partial, and normal. The type of restore that you select will depend on the type of backup that was performed. You can initiate each restore type from the NetWorker User for SQL Server GUI Windows installed on the database server.

The nsrsqlrc program, installed with the NMSQL software, performs the restore process. The API translates the object names requested for restore by nsrsqlrc into a format NetWorker recognizes so that the data can be requested to be received from the storage node nsrmmd process.

You must enable the SQL Server filestream feature on the destination recovery server to restore SQL Server 2008 filestream data.

NetWorker Module for Microsoft Exchange

The NetWorker Module for Exchange (NME) can be used to link your NetWorker backup and recovery environment with Microsoft's Exchange Server to protect the Exchange mail server.

The NetWorker Module for Microsoft Exchange provides some additional programs on the database server to perform online, and transaction log backups of the Exchange application. It also provides the special NetWorker User for Exchange GUI for manual backup and restore operations. NME can perform online mail server backups including the Information Store, Directory Store, individual mailboxes, and transaction logs.

Backups with NME 5.1 can perform backup and restore through the traditional Exchange aware Backup API or using the Backup Utility Exchange agent. With the traditional backup method, NME will contact the Exchange database through a local command and request a consistent backup through the supported Exchange Backup API.

The Backup Utility Exchange agent can be used for mailbox level (brick-level) and public folder backup and restore operations. The Backup Utility Exchange agent uses MAPI to connect to each mailbox and public folder the same way Outlook does (instead of using the API). Download and install the Microsoft Exchange Server MAPI Client and Collaboration Data Objects kit from Microsoft to have support for mailbox and public folder backup options for NME 5.1 running on Windows 2008.

Snap-shot type backups with Microsoft Volume Shadow Copy Service (VSS) or NetWorker PowerSnap module are not currently supported by NME version 5.1. You can achieve Snapshot VSS backups of Exchange through the NetWorker Module for Microsoft Applications.

Which Server are you Backing Up?

Client resource

When performing NME backups, create multiple Client resources using the Exchange server hostname to perform different types of scheduled backups at different times. These client resources can start full information store backups, transaction log backups and file system backup data. You could also create optional clients for mailbox and public folder backups.

When Should the System be Backed Up?

Group resource

You can start an NME backup as a manual or scheduled backup. With a manual backup, the administrator can choose different components of the Exchange server to backup through the NetWorker User for Exchange Server GUI.

It might be sufficient to backup the Exchange server as a Full backup once a week with a transaction log when creating groups for a smaller Exchange environment. You could then perform differential or incremental backups in the middle of the week. You should perform transaction log backups more frequently than full database backups to bridge the gap of time from one full backup to another.

What Level of Backup Should be Performed?

Schedule resource

You can request five Exchange Backup levels through a configured schedule resource. These are full backup, copy backup, incremental backup, differential backup and skip a backup. These options match the NetWorker Schedule resource in the following chart:

Exchange Backup level	NetWorker Schedule	What is backed up?
Full	Full	Full database, then transaction logs. Logs truncated after successful backup.
Differential	1 and 2-8	Transaction logs backed up. Logs not truncated after successful backup so they are there for a following backup. Always recorded as level 1.
Copy	9	Database and transaction logs copied. Logs not truncated after successful backup.
Incremental	Incr	Transaction logs backed up. Logs truncated after successful backup so they are not there for a following backup.
Skip	Skip	Nothing is backed up– don't perform the backup

In addition, it is important to note that if you are planning to take advantage of the mailbox and public group backups there are some variations on the supported NetWorker levels. They are documented here:

NetWorker Schedule	Exchange Level	Mailbox or Public Folder backups
Full	Full	Each object is backed up
Levels 1-8	Differential	Backup will be Incremental
Level 9	Copy	Backup will be level Full
Incr	Incremental	Items modified or created since the previous backup (full or incremental) are backed up
Skip	Skip	Nothing is backed up– don't perform the backup

What Do you Want to Back Up on That System?

Save set attribute in Client resource

The NetWorker client used for your NME backups will need to have a special save set entry to indicate an Exchange save set is being backed up. The following chart shows the save set values and a description of what each will back up:

Save set attribute	What is backed up?
MSEXCH:	Full Exchange database, including Information store and directory store using traditional backup API
MSEXCH: IS	Information Store only
MSEXCH:IS/Storage Group1	Backup only StorageGroup1 from Information Store
MSEXCH:IS/Storage Group2	Backup only StorageGroup2 from Information Store
MSEXCH:PF	Public Folder backups
MSEXCH:MB	Mailbox level backups
MSEXCH:MB/Username	Mailbox level backup of "Username's" mailbox only
MSEXCH:MB/Username/Inbox	Mailbox level backup of "Username's" Inbox only
MSEXCH:MB/Us*	Mailbox level backup of all names beginning with "Us"

The mailbox level backups are useful for recovering individual mailboxes when a specific user loses email, and that is the only mailbox that must be restored. Use this strategy for a targeted number of individuals since this method of backup is slow.

What Command Should be Used to Perform the Backup?

Backup Command attribute in Client resource

The Backup command attribute in the client resource is used to call the custom NME programs nsrxchsv or nsrxchmsv. They are used to interface with the Exchange backup API and Exchange MAPI backup agent.

The nsrxchsv program interfaces with the Exchange backup API to request the traditional information store, mailbox and public folder backups of Exchange 2000 and Exchange 2003. For Exchange 2007, the nsrxchsv command is only used for Information Store and public folder backups. The nsrxchmsv program interfaces with the Exchange 2007 MAPI backup agent to perform mailbox level backups.

If you are backing up Exchange running on a virtual server, you will also want to specify the “-a virtual_server_name” option used to tell NetWorker the client is a virtual server.

Microsoft Exchange Restore

You can perform restore operations from the NetWorker User for Exchange Server with NME, the same interface used for manual backups. The options that are available for restore will depend on the type of backup that was performed. The following table illustrates some of the restore types that are available when appropriate backup data and media is available.

Object backed up	Object Recoverable
Information Store	<ul style="list-style-type: none">- Entire Information Store- One or more storage groups- Public folder database or mailbox database.
Storage group	<ul style="list-style-type: none">- Entire storage group- One or more databases in the storage group
All public folders	<ul style="list-style-type: none">- Each public folder tree- One or more individual public folders- Individual items in a public folder
A public folder tree	<ul style="list-style-type: none">- Entire public folder tree- One or more public folders in the tree- Individual items in a public folder
A single public folder	<ul style="list-style-type: none">- Public folder- Individual items in the public folder

Object backed up	Object Recoverable
All private mailboxes	<ul style="list-style-type: none"> - One or more mailboxes - One or more folders in a mailbox - Individual items in a mailbox
A single mailbox	<ul style="list-style-type: none"> - Mailbox - Individual items in the mailbox
One or more individual mailbox or public folder items	<ul style="list-style-type: none"> - Individual items

The nsrsqlrc program installed with the NME software performs the restore process. The Exchange API translates the object names requested for restore by nsrxchrc into a format that NetWorker recognizes when restoring data backed up through a traditional method. The nsrxchmbrc command is used instead of the normal nsrxchrc command when performing a restore of a mailbox level object or item for Exchange 2007.

When restoring a storage group on your Exchange server (2003, 2007), you can choose to restore to a Recovery Storage Group (RSG). This allows you to recover a storage group to an alternate location within the Exchange server (the RSG) and then choose relevant information from the recovered copy.

When recovering the entire information store or a storage group, the NetWorker User for Exchange Server interface gives an option to also “Complete Exchange recover process” (replay logs). It puts the database online after a successful restore/recover process is completed. If complete replay of logs is not performed by the NetWorker restore process, it will need to be done afterwards using the eseutil Microsoft Exchange Server utility.

NetWorker Module for Microsoft Applications

The NetWorker Module for Microsoft Applications (NMM) can be used to link your NetWorker backup and recovery environment with Microsoft’s Volume Shadow Copy Services (VSS) to protect different Microsoft applications and operating system components.

We will not cover the NMM in detail in this article since EMC already has an NMM whitepaper, but a brief explanation is given here so that you understand where NMM fits in comparison to the other NetWorker Module solutions for SQL and Exchange.

VSS was first introduced in Windows 2003 and is now part of Windows 2008. The VSS framework has three major components: writers, providers and requestors. The writer is the component of the application being backed up that guarantees there is a consistent copy of the data to backup. The provider is the storage component that keeps the different point-in-time copies (snapshots) of the data, and the requestor is the software that initiates the creation of shadow copies and performs the backup.

The NetWorker software requests a Shadow Copy snapshot be created with the NMM. For every application or component that has a writer, NetWorker can request creation of a snapshot to capture the data for backup. NetWorker uses the snapshot to create a single consistent point-in-time copy of the application that can then be copied to traditional backup media such as tape, virtual tape or disk.

Some of the applications that can be backed up and protected by NMM include SQL Server, Exchange Server, SharePoint Server, Data Protection Manager, Active Directory, Hyper-V, and File system snapshots.

The NMSQL and NME Modules use more traditional methods and APIs for backing up the databases and do provide some features that are not supported with NMM such as mailbox level backup and restore.

There is no question there is some overlap between NMSQL, NME and the NMM, but the approach for backup and recovery is different (API vs. VSS). There is a very good chance that NMM and VSS backup will be the future of Microsoft Application backup and restore as more features and applications are added, but for now we have more than one option available for SQL and Exchange backups.

Please check powerlink.emc.com for more information on NMM such as the “NetWorker Module for Microsoft Applications 2.1 Administration Guide” and the “Backup and Recover Using the EMC NetWorker Module for Applications” whitepaper.

NetWorker Module for SAP R/3

The NetWorker Module for SAP R/3 with Oracle (NMSAP) can be used to link your NetWorker backup and recovery environment with SAP R/3 running on Oracle through the SAP R/3 BR*Tools program. The main component of the BR*Tools program is the backint program which provides access for between SAP R/3 (SAP) and NetWorker to perform backup, inquiry and restore processes for all Oracle and SAP files.

The NetWorker Module for SAP R/3 provides additional programs on the SAP server used to perform online, offline and archived redo log backups. NMSAP can also be used in conjunction with the sapclone feature of BR*Tools to make additional copies of the backup data for offsite protection, or to copy backup data between disk and tape media. The main supporting programs for SAP BR*Tools that interface with backint and NetWorker are brbackup, brarchive, brrecover and brrestore. You can also use an SAP BrGui and character based “brtools” program to interface and call the other BR*Tools commands.

Which Server are you Backing Up?

Client resource

Create multiple client resources using the SAP server hostname to perform different types of scheduled backups when configuring NMSAP. Examples of the different backup types include full online database backups, archive redo log backups and backing up file system data such as the backint log files or other operating system data.

When Should the System be Backed Up?

Group resource

You can start an NMSAP backup as a manual or scheduled backup. You must perform three configuration tasks before this can happen.

Before the NMSAP backint can be used for any backup and restore operations the first task is to copy the backint (UNIX/Linux) or backint.exe (Windows), which is part of the NMSAP software to the same location as the BR*Tools programs.

The second task is to update the SAP initialization file. The SAP initialization file is normally named initSID.sap where SID is the Oracle system identifier (SID) and is typically located in \$ORACLE_HOME/dbs on UNIX and Linux systems and %ORACLE_HOME%\database directory on Windows.

The backup device type parameter and the backup utility parameter file are two required settings that need to be updated in the initialization file.

For the backup device type specify either:

```
backup_dev_type = util_file
```

or

```
backup_dev_type = util_file_online
```

For the backup utility parameter file, update it to point to the NMSAP utl file:

```
util_par_file = $ORACLE_HOME/dbs/initSID.utl
```

or

```
util_par_file = %ORACLE_HOME%\database\initSID.utl
```

The third task is to create a proper parameter file for NMSAP. When performing a manual or scheduled backup, the backint and BR*Tools programs use parameters in the parameter file, initSID.utl. Be sure to have the initSID.utl file properly configured before attempting any backint backup. Examples of some parameters that can be set include client, server, pool, group, etc. Most of the parameters are optional or should be left at

default values. If updated, they will take precedence over any NMC GUI settings. There is a full list of the parameters and their descriptions in the NetWorker Module for SAP with Oracle Multiplatform Version Administrator Guide.

With a manual backup, the administrator can use the BrGui interface to begin a backup from the SAP server or run a BR*Tools command (brbackup/brarchive) directly from the command line or a script. When performing a manual brbackup with backint it will send Oracle datafiles, control files, existing profiles and SAP internal catalogs as backup sessions to NetWorker. When using brarchive the archived redo logs will be passed for backup to NetWorker.

For smaller SAP environments, it might be sufficient to backup the server as a Full backup once a week with archived redo logs backing up in the middle of the week. You should perform archived redo log backups more frequently than the full database backups and bridge the gap of time between the full database backups.

When backing up SAP on Windows in a scheduled backup, you will need to update the Remote User and Password attributes in the SAP client resource for proper authentication to the SAP application.

When SAP is running with Oracle, the RMAN utility is also an option for capturing database data and archived redo logs for backup. However, using RMAN in this scenario is optional since the NMSAP backint program can be used to capture all necessary data from the database. Although RMAN can be used for performing the backup, recovery and restore operations, this article focuses strictly on using the backint for performing these operations.

What Level of Backup should be Performed?

Schedule resource or configuration file

When configuring a schedule for your NMSAP backups, the level of backup can be either a full database backup with brbackup or an archived redo log backup with brarchive. The level of backup is not set in the Schedule resource as with most NetWorker client backups. Instead, the NMSAP configuration file can be used to decide what type of backup to perform. Do not select a level Skip backup in an assigned schedule for a day that you do want to perform a backup.

What Do you Want to Back Up on that System?

Save set attribute in Client resource

The NetWorker client used for your NMSAP backups will need to have a save set entry to indicate an SAP backup is taking place and which Oracle SID to back up. Use the keyword "backint:" and follow it with the Oracle SID to indicate the SAP backup is to take place.

Save set: backint:SID

For NMSAP, we use the configuration file (nsrsapsv.cfg) covered in the next section to determine the actual data to backup.

What Command should be Used to Perform the Backup?

Backup Command attribute in Client resource

The Backup command attribute in the client resource calls the custom NMSAP program `nsrsapsv` that will be used to interface with `backint` to request the backup. The `nsrsapsv` command is put in the backup command attribute with a “-f” option to specify the location of the configuration file with `backint` backup settings. The format of the setting will be: `nsrsapsv -f filepath/nsrsapsv.cfg` where `filepath` is the complete path to the file on the NetWorker client and `nsrsapsv.cfg` is the name of the NMSAP configuration file.

A template of the NMSAP configuration file, `nsrsapsv.cfg` is installed with the NMSAP client software. On UNIX and Linux, it will be located in `/etc`; on Windows you will find it in `C:\Program Files\Legato\nsr\res`. Copy this template file and update it to contain proper settings for your setup and specific Oracle SID. You will likely have multiple `nsrsapsv.cfg` files for different types of backups. For example, you might have an `nsrsapsv_SID_full_database.cfg` file to call `brbackup` to perform a full database backup and an `nsrsapsv_SID_archive.cfg` for calling `brarchive` to perform archived redo log backups.

This chart describes the parameters that require updating:

Required Parameter	Description
BR_EXEC	Choose <code>brbackup</code> (full backup) or <code>brarchive</code> (log backup)
HOMEDRIVE	Drive letter where Windows system files are located – not used for SAP on Unix and Linux.
ORACLE_HOME	Name of operating system user that runs the Domino server
SAP_BIN	Directory path to location of BR*Tools binaries and <code>backint</code> program file included with NMSAP.
SAPBACKUP	Directory to store backup logs. Required for Windows. Unix and Linux will use <code>\$ORACLE_HOME/sapbackup</code> as a default location.
ORACLE_USR_PASSWD	Oracle username and password
OS_USR_PASSWD	SAP operating system username and password

Use the nsrsapadm command to update the ORACLE_USR_PASSWD and OS_USR_PASSWD parameters in the configuration file. This command will encrypt the username and password so that it cannot be read when viewing the contents of the file.

Command syntax:

```
# nsrsapadm -c nsrsapsv_custom_file.cfg
```

where nsrsapsv_custom_file.cfg is the configuration file you want updated.

Backup session reports can be a useful tool for monitoring or troubleshooting your SAP manual and scheduled backups. The backup session reports can be found in:

C:\Program Files\Legato\applogs\nsrsapsv.<process_id>SID for Windows

or for UNIX and Linux in: /nsr/applogs/ nsrsapsv.<process_id>SID

You can also check the backintSID.log.raw found in the same location.

SAP R/3 with Oracle Restore and Recovery

SAP R/3 data backed up with NMSAP and the brbackup, brarchive BR*Tools commands can be restored from NMSAP by using either the BrGui interface supplied by SAP R/3, the brtools character-based program or the brrestore, brrecover command directly.

Verify that both the SAP initialization file, initSID.sap and the initSID.utl file are configured correctly before attempting a restore or recover operation.

Verify that a required backup exists before initiating the restore or recovery operation. Perform this by using the “-verify” option on the brrestore command.

After a restore is complete, you can find results in the backintSID.log.raw file in the applogs sub-directory wherever the NetWorker nsr directory is located on the client.

Conclusion

Due to costs or other organizational constraints, system administrators are being asked to take on more tasks. Performing backups is only a portion of their job roles.

Researching the required NetWorker Modules in great detail for each database is ideal, but requesting outside help from a services organization such as EMC Global Services would ensure that you get the best solution. Sometimes, these options are just not possible. This article was written to give guidelines on the similar properties of the NetWorker Modules and a more detailed look at NMO, NMSQL, NME, NMM and NMSAP when an administrator might need to deploy the backups themselves. It is also a useful reference after an external services organization has configured the backups.

Every application has transactional data that needs to be captured through backups. We have covered some of the strategies to perform these backups. Deployments of NetWorker Modules not covered in detail in this article can also benefit from some of the information covered on configuring clients, backup times, backup levels, retentions, pools, devices and directives.

Always remember that backups are being done to address situations when a restore or recovery operation is needed. Consider the restore and recovery SLAs of a particular database and make sure these requirements can be met. Information in this article can help to build and configure your backup strategy, but also remember to regularly test your application restores.

References

1. EMC NetWorker Module for Oracle 4.5 Multiplatform Version Administration Guide, EMC Corporation, May 2007
2. EMC NetWorker Module for Microsoft SQL Server Release 5.2 Administration Guide, EMC Corporation, September 2008
3. EMC NetWorker Module for Microsoft Exchange Server Release 5.1 Administration Guide, EMC Corporation, August 2007
4. EMC NetWorker Module for Microsoft Applications Release 2.1 Administration Guide, EMC Corporation, September 2008
5. EMC NetWorker Module for SAP with Oracle Release 3.5 Multiplatform Version Administration Guide, EMC Corporation, August 2007