

# SAP Backup Solutions using EMC NetWorker

*Applied Technology*

---

## **Abstract**

The purpose of this technical paper is to present methodologies for backing up SAP databases using EMC<sup>®</sup> NetWorker<sup>™</sup> capabilities while exploiting EMC storage technologies. Modules for SAP using BRBACKUP for Oracle, Microsoft SQL Server, IBM DB2 Universal Database, and IBM Informix are introduced. Also included is an introduction to NetWorker disk backup options for SAP customers.

May 2006

---

---

Copyright © 2006 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.

All other trademarks used herein are the property of their respective owners.

Part Number H2140

---

## Table of Contents

<b>Executive summary</b> .....	<b>4</b>
<b>Introduction</b> .....	<b>4</b>
<b>Audience</b> .....	<b>4</b>
<b>Terminology</b> .....	<b>4</b>
<b>Overview of NetWorker architecture</b> .....	<b>5</b>
NetWorker server .....	6
NetWorker storage node .....	6
NetWorker client .....	7
<b>Best practices</b> .....	<b>7</b>
LAN-free storage node .....	7
<b>NetWorker capabilities</b> .....	<b>9</b>
NetWorker PowerSnap .....	9
NetWorker Module for SAP R/3 with Oracle (NMSAP) .....	10
Backup with SAPDBA or BRBACKUP .....	10
Restore with SAPDBA or BRRESTORE .....	11
Configuration to use NMSAP Module .....	11
NetWorker Module for SQL (NMSQL) .....	12
NetWorker software .....	13
NetWorker Module for DB2 Universal Database (NMDB2) .....	14
NetWorker Module backup .....	14
NetWorker Module restore .....	14
The db2uext2 program and log files .....	14
NetWorker Module for Informix (NMINF) .....	15
NetWorker Module for Informix backup .....	15
NetWorker Module for Informix recovery .....	16
<b>Split-mirror backup with SAP databases and NetWorker</b> .....	<b>17</b>
The split-mirror concept .....	17
Method of backup with split mirror .....	17
Options for creating the split-mirror replica .....	17
Split mirror with BRBACKUP for Oracle .....	18
<b>New options for backup and recovery</b> .....	<b>18</b>
Considerations when selecting backup-to-disk options .....	18
Advantages of backup to disk .....	19
Disk-based backup: Two enterprise options .....	19
Choosing a disk-based backup option .....	20
EMC recommendations for B2D .....	21
A two-tiered backup approach (D2D2T) .....	21
<b>Conclusion</b> .....	<b>22</b>
<b>References</b> .....	<b>22</b>

---

## Executive summary

Using EMC® NetWorker™ to protect critical SAP data provides flexibility when architecting the best technical and business solution for backup and recovery of critical SAP data. NetWorker protects your data by simplifying and centralizing backup and recovery operations.

NetWorker solutions are described for backup and recovery of each of the databases found in SAP landscapes. Included is a discussion of SAP and split-mirror backup. A summary of backup-to-disk (B2D) technologies completes the backup solution options presented. Best practices and EMC architectural recommendations are provided where applicable.

## Introduction

As a general rule, the backup strategy for an SAP system landscape depends on its recovery requirements. Recovery scenarios for different error situations may require specific backup procedures. Apart from application data, the software and configuration files of the system components must also be included in a backup and recovery plan.

When designing a backup strategy for SAP landscapes, different kinds of data must be protected:

- Data managed by databases
- Data in files which are managed directly by the applications
- Application software
- Configuration files for software components
- Log files of software components

Each of the data types may use a different backup method. File system data, which covers most of the previous list, should be backed up regularly with frequency depending upon the importance of interface and log files—or as changes are made to application software and configuration files. SAP best practice recommendations are database backup daily, with database log backups occurring multiple times during each day.

This paper focuses on data managed by databases. The goal of this paper is to present EMC options and technologies available for backup of SAP landscapes. NetWorker's versatility and flexibility allow several possible configurations for LAN-based and LAN-free backup. All enterprises have some form of unique requirements or configurations along with their SAP landscape. The intention of this paper is to present some guidelines and introduce backup offerings for data protection within SAP landscapes.

## Audience

The audience for this paper is customers, including IT planners, SAP Basis administrators, SAP database administrators, storage architects, administrators, and any others involved in evaluating, acquiring, managing, operating, or designing an EMC networked storage environment. This paper is also aimed at EMC staff and partners for guidance and development of proposals.

## Terminology

**client file index (CFI)** – NetWorker database(s) that tracks each file (pathname) in a client's backup, allowing clients to browse their backups for files from a particular point in time. The NetWorker server creates and maintains one client file index per client.

**media database** – NetWorker database used to track the backups and the volumes that store the backups.

---

**NetWorker control data** – Collectively represents the NetWorker configuration information and the backup tracking information stored on the NetWorker server. NetWorker control data includes: resource directory, media database, client file index.

**nsrd** – Network save & recover daemon. Master daemon. Provides NetWorker server services; reads and maintains resource files for configuration information. Also starts the nsrmmdbd and nsrindexd daemons. Also starts the nsrmm(s) for device(s) local to the NetWorker server.

**nsrexecd** – Network save & recover execution daemon. Runs on NetWorker clients to support remote execution requests from NetWorker servers. Also determines which ports to use to support and request NetWorker services.

**nsrindexd** – Network save & recover index daemon. Provides read and write service for the client file index (CFI) databases on the NetWorker server. Started by nsrd.

**nsrmm** – Network save & recover media multiplexing daemon. The NetWorker storage node daemon. Runs on the storage node(s) and supports reading and writing backup data to devices and volumes. One nsrmm executes for each configured device.

**nsrmmdbd** – Network save & recover media management database daemon. Started by nsrd. Supports read and write services for the media database.

**resource directory** – Contains the configuration resource files.

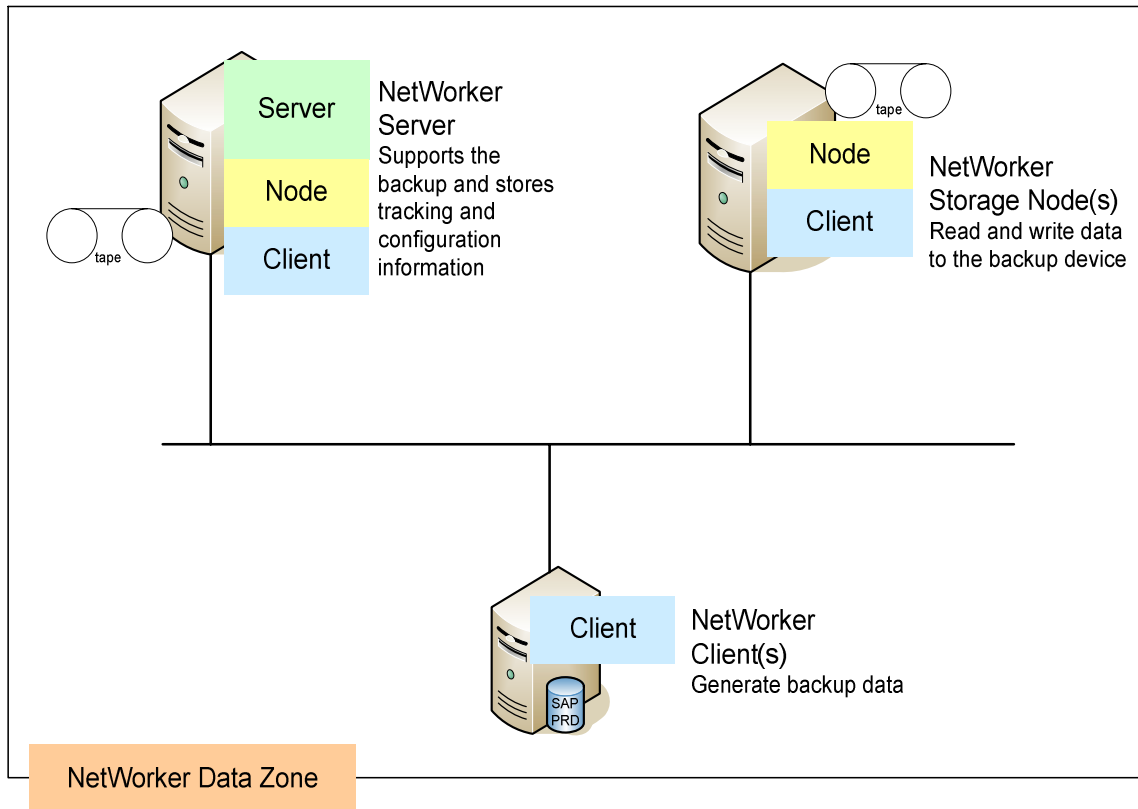
**resource files** – Stored and maintained configuration information of defined resources; where each resource represents a configurable NetWorker component.

## Overview of NetWorker architecture

NetWorker is based on client/server architecture. A NetWorker datazone consists of a single NetWorker server, NetWorker clients, and NetWorker storage nodes. Figure 1 on page 6 shows an example of a NetWorker datazone.

Typically, a NetWorker server hosts one or more backup devices (tape or disk). When a backup occurs, a NetWorker client sends data across the LAN to the NetWorker server or storage node, which directs the data to appropriate storage. The server receives metadata information about the files being backed up separately and makes appropriate entries in the client file index and media database.

NetWorker manages client backup data in increments called *save sets*. A save set originates from a single client and may consist of a single file, a directory, a file system, or a partition. The NetWorker client file index on the server contains detailed information about each save set (for example, the filename of each file stored, time of backup, size, permissions, ownership, and so on). An entry for each save set is also contained in the server's media database. The media database maps the save sets to the volumes where they are stored, stores the browse and retention policies for each save set, and maintains tracking information for all storage volumes. The client file index and media databases are essential to NetWorker's ability to locate and recover data rapidly.



**Figure 1. NetWorker client/server architecture**

## ***NetWorker server***

The NetWorker server is the datazone host that stores the configuration information, such as supported clients, devices, when to run the backups, and so on. The NetWorker server also stores the online NetWorker databases that track the backups and volumes. The NetWorker server, as a client of itself, automatically backs up the configuration and tracking databases to protect NetWorker data.

There is a single NetWorker server per datazone, and it must be available for any NetWorker activity to be supported in that data zone. NetWorker servers have NetWorker client, storage node, and server software installed.

## ***NetWorker storage node***

NetWorker storage nodes are datazone hosts with directly attached (or SAN-accessible) devices that provide the NetWorker interface for the backup devices and volumes. If devices are not accessed through the NetWorker server, they are considered remote devices, and the host through which the devices are accessed is a remote storage node.

The NetWorker server is always a storage node in any datazone and is the default storage node for backups. Using remote storage nodes is optional, although they distribute the backup workload and reduce network traffic. Storage node hosts in the datazone have both the NetWorker client and storage node software installed.

NetWorker storage nodes substantially increase the parallelism and resiliency of a NetWorker configuration, and have proven highly useful in many NetWorker installations. A storage node is connected to one or more storage devices. The NetWorker server may direct a client to send its data to a storage node for backup. It is also possible to have NetWorker client and storage node software both installed on the same system. Although client data is sent to the storage node, client metadata is always

---

directed to the NetWorker server. The server maintains the file index and media database information for all data on the storage node.

In a backup, NetWorker clients send backup data to a particular storage node based on each client's configuration. During recovery, the client reads from the storage node that provides the device that contains the necessary volume.

Storage nodes also send tracking information about the volume and any data written to it to the NetWorker server to store it as online tracking information about the volume and its contents.

## **NetWorker client**

The NetWorker client is the largest NetWorker software component and the fundamental host. The client's most important functions are to generate backups (called save sets), push them to the storage node, and retrieve them during a recovery. NetWorker clients in an SAP landscape are the database servers and application servers.

A NetWorker client is any system with data to be backed up—whether a desktop system or a large database server. NetWorker client runs client software tailored to the operating system (and, optionally, the database and messaging product) it uses. Normally, a client operates under the control of the NetWorker server according to guidelines and to schedules an administrator establishes. An administrator can also configure NetWorker to allow a user to initiate manual backups and restores on the client system as necessary. For instance, if a particular file must be recovered from backup, that action can be initiated from the client without administrator intervention. Likewise, the user can manually back up important files without waiting for regularly scheduled backups to occur.

While performing a backup, the client also generates tracking information (file and directory names included in the backup and the time of the backup) and sends it to the server for use in point-in-time recoveries.

Additionally, the client provides the user backup, recover, and administrative interfaces to allow configuration and monitoring of NetWorker for your specific environment. NetWorker client software is installed on all participating hosts in the datazone, including hosts that also perform the roles of server and storage node.

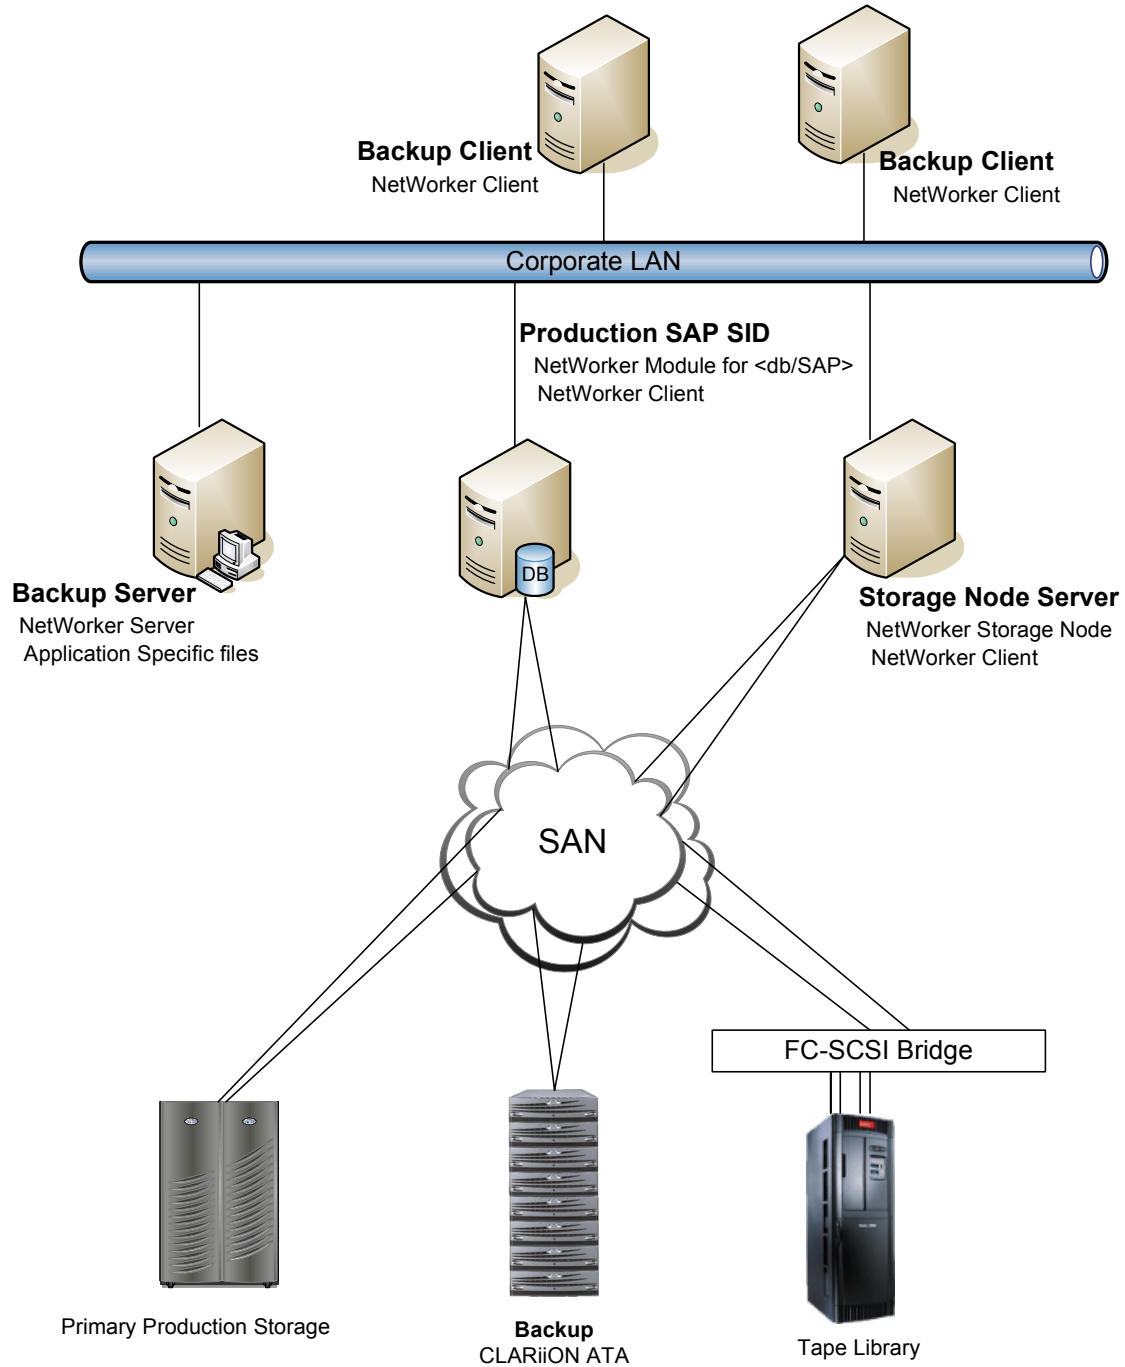
## **Best practices**

The following example outlines best practices and configuration guidelines for protecting SAP data. This architecture design is included to assist in backup and recovery infrastructure decision making in your enterprise environment. Backing up and restoring typically involves a large quantity of data to be moved to and restored from disks. As a best practice, one should strive for a LAN-free approach to minimize the impact to LAN users and applications. As every environment is unique in its requirements and configurations, the intention of this paper is to present base guidelines for data protection. For more information tailored to best suit your SAP environment, contact your local EMC Sales Representative. You can learn read more about NetWorker from at EMC's Powerlink™ support site at <http://Powerlink.EMC.com>.

## **LAN-free storage node**

A large file server or database server may be configured as a storage node primarily for the backup of local data. Such a configuration prevents data from traversing the IP network, maximizing backup and recovery performance, and reducing network congestion. Storage nodes provide this capability within SAN environments. A storage node does not provide service to network clients. Instead, a storage node is configured to back up its data directly to a SAN-attached storage device, such as a tape library or disk. Multiple dedicated storage nodes can share a single storage device (such as a tape library or CLARiiON storage system).

In a LAN-free backup, NetWorker sends backup traffic over the Fibre Channel network or direct-attached network, while a small amount of the backup client's metadata traffic and cataloging information is sent to the NetWorker server through the LAN as shown in Figure 2 on page 8. One option is to install the storage node and the NetWorker client on the same host. This is referred to as a dedicated storage node. A possible disadvantage to this configuration is a limitation of what can be backed up (dependent on the individual environment). A dedicated storage node cannot backup other clients in the network.



**Figure 2. Example architecture for LAN-free storage node**

---

It is not likely that every single client will have the storage-node software. In a large environment with many clients, a small number of storage nodes will typically serve as an aggregation point for a large number of clients. The clients will send their data to their associated storage node; the storage node will store the data on the backup device.

When maximum capacity of a storage node is near, additional storage nodes can be installed on the network to alleviate overloading of a single storage node. Storage nodes can share a single backup storage device, EMC CLARiiON® storage system, or tape library.

## NetWorker capabilities

There are many pieces in an overall solution for your SAP landscape and any additional business requirements must also be considered.

EMC NetWorker is a high-performance backup and recovery software solution. NetWorker enables SAP customers to standardize on one application to provide complete, reliable protection in the largest data centers and the smallest satellite offices. The NetWorker software is built upon an open, highly scalable architecture, providing data protection capabilities and management consistency required by heterogeneous enterprises.

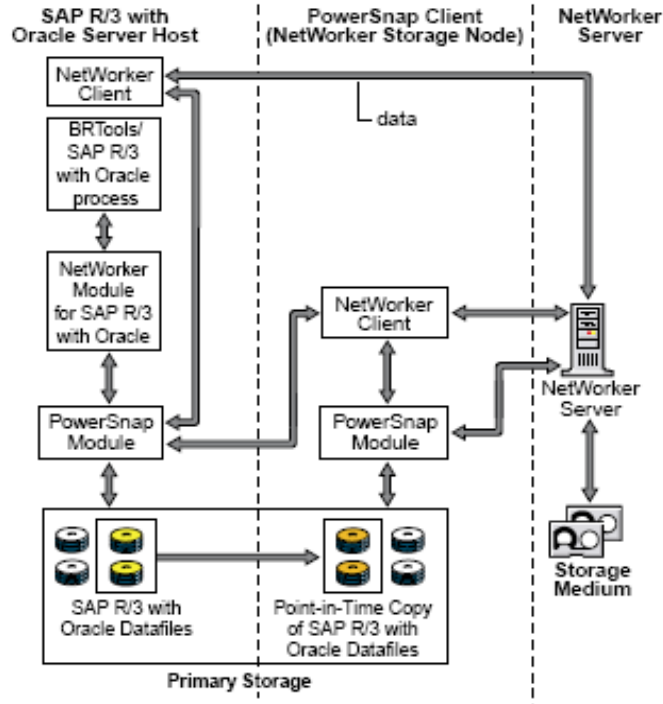
NetWorker application modules provide automated, online protection and granular recovery (datafiles, tablespaces, messages) for leading databases, email, and SAP applications. These modules can eliminate the need for custom scripting and through full integration with vendor APIs, enable SAP customers to achieve 24x7 application access, ensure application integrity, and provide reliable protection.

## NetWorker PowerSnap

NetWorker PowerSnap™ Module is available on EMC Symmetrix® and CLARiiON storage systems, as well as other storage arrays. NetWorker PowerSnap works with the NetWorker server and NetWorker Module for SAP to perform backups and restores of data residing on the specific types of primary storage. The backup client can leverage PowerSnap functionality to instantly create the point-in-time (PIT) copy of the snap set on the primary storage volume as a snapshot session. The PIT copy is registered with the NetWorker server. A restore of snapshots where saved data is retrieved from a mounted PIT copy or from disk device can occur instantaneously.

To use the PowerSnap features, the NetWorker Module must be used with the EMC NetWorker PowerSnap Module. The PowerSnap Module creates the snapshots, and the NetWorker Module stores the snapshot data. The NetWorker Module for SAP (NMSAP) implements SAP **backint** specifications for backup and restore purposes.

Figure 3 on page 10 shows backup and recovery data flow with NetWorker PowerSnap.



**Figure 3. NetWorker Module PowerSnap backup and restore data flow**


### ***NetWorker Module for SAP R/3 with Oracle (NMSAP)***

The NetWorker Module for SAP R/3 with Oracle software is an add-on module for NetWorker client software that provides a backup and restore interface between the SAP BRTools program and the NetWorker server. The main component of the NetWorker Module is the **backint** program, which provides backup, inquiry, and restore processes for all Oracle and SAP R/3 files.

#### **Backup with SAPDBA or BRBACKUP**

When a backup is initiated from **SAPDBA** or the **BRBACKUP** program on the SAP R/3 server running Oracle, **BRBACKUP** sends a request to the **backint** program. The **backint** program runs two or more times as required to back up the requested Oracle datafiles, control files, online redo logs, and profiles. It sends all the data and process information in parallel streams to the NetWorker server. The datastream is backed up on storage media or a NetWorker storage node, and NetWorker software creates entries in the client file index and media database. The final results of the backup are recorded in a log file named `backint<ORACLE_SID>.log` on the SAP R/3 server running Oracle.

Figure 4 on page 11 shows how the NetWorker services, processes, and programs interact during a backup.

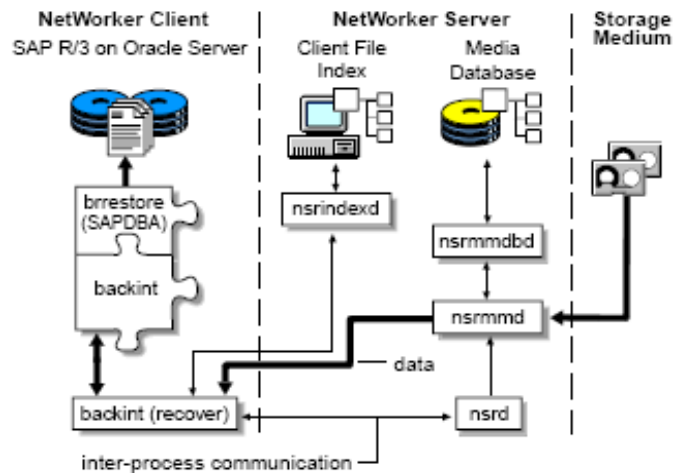


**Figure 4. NetWorker and database interaction during a backup**

### Restore with SAPDBA or BRRESTORE

When a restore is initiated from **SAPDBA** or the **BRRESTORE** program on the SAP R/3 server running Oracle, **BRRESTORE** sends a restore request to the **backint** program. The **backint** program browses the NetWorker server's client file index and restores the specified files in a parallel stream from the NetWorker server to the SAP R/3 server running Oracle.

Figure 5 on page 11 shows how NetWorker services, processes, and programs interact during a restore.



**Figure 5. NetWorker and database interaction during a restore**

### Configuration to use NMSAP Module

The NMSAP module can perform backups using operating system tools such as `dd` or `cpio`. To use operating system tools, configure the SAP initialization file `init<SID>.sap`, typically located in `%ORACLE_HOME%\DATABASE`, to point to the **backint** parameter file which is called `init<Oracle_SID>.utl`.

---

Adjust the following parameters in the `init<SID>.sap` configuration file:

- `backup_dev_type=util_file`
- `util_par_file=?/DATABASE/init<Oracle_SID>.utl`

Or for Windows:

- `util_par_file=?\database\init<Oracle_SID>.utl`  
(where ? = %ORACLE\_HOME%)

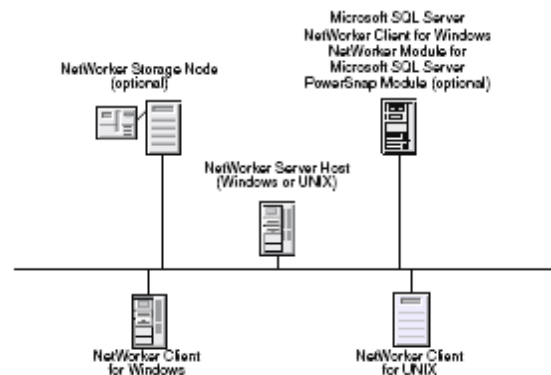
Adjust the following parameters in the `init<Oracle_SID>.utl` configuration file:

- `group=Default`
- `pool=Default or any pool defined in your NetWorker server`
- `expiration=Time in days, weeks, months`
- `server=NetWorker server hostname for backups and restores`

## NetWorker Module for SQL (NMSQL)

The NetWorker Module for Microsoft SQL Server is an add-on module for NetWorker software that provides back up and restore of Microsoft SQL Server 7.0 and SQL Server 2000 databases and transaction logs. The NetWorker Module integrates data protection procedures for Microsoft SQL Server databases with the NetWorker software. This provides a comprehensive storage management solution that addresses the need for cross-platform support of enterprise applications on UNIX and Microsoft Windows operating systems.

Typically, the NetWorker server software resides on one server and the SQL Server database and software on another. Figure 6 on page 12 shows an example configuration.



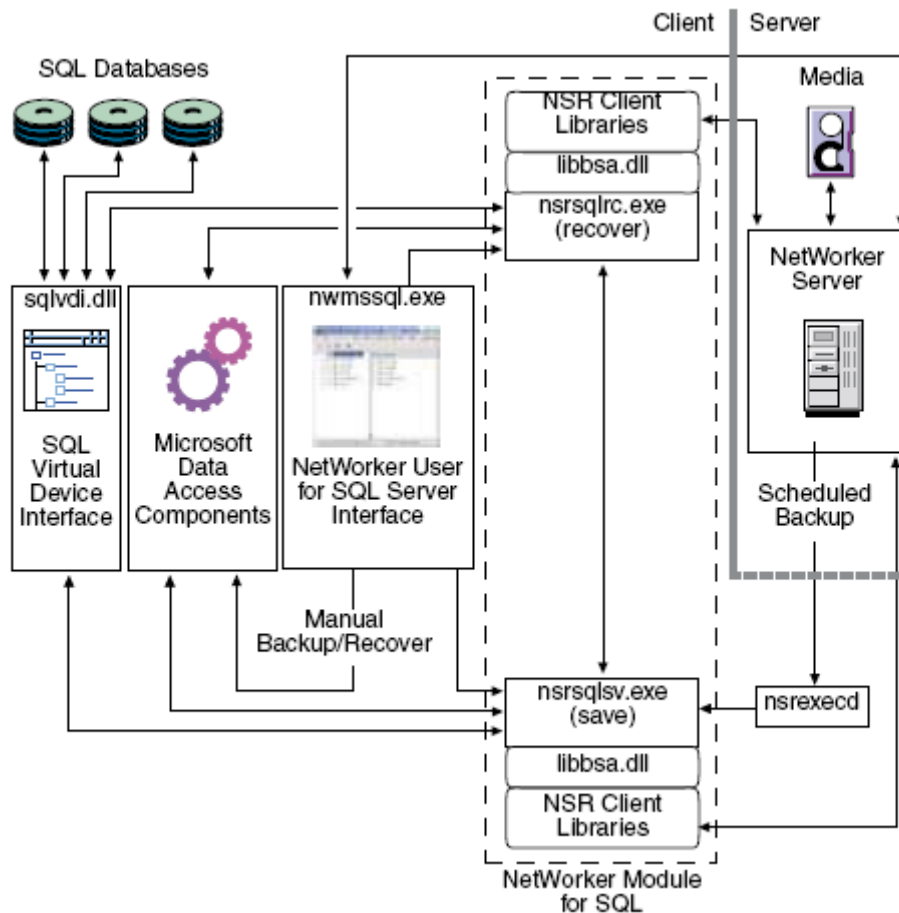
**Figure 6. NetWorker with Microsoft SQL Server Module example architecture**

The NetWorker client software and the NetWorker Module for Microsoft SQL Server software must both be installed. The NetWorker server software may be installed on the SQL server host or on a separate computer.

- A NetWorker server is a system running the NetWorker service software with a connected backup storage device.
- A NetWorker client is a system running the NetWorker client software, which receives storage management services from the NetWorker server.

After installation of the NetWorker Module, you must configure the SQL Server host as a client of the NetWorker server.

Figure 7 on page 13 shows how the NetWorker Module and NetWorker software work with SQL Server.



**Figure 7. NetWorker Module and NetWorker setup with SQL Server**

## NetWorker software

The NetWorker client software must be installed on the SQL Server host computer. The *NetWorker Installation Guide* provides more information about the installation requirements. The NetWorker Module can connect to and interact with a NetWorker server running on any supported operating system platform.

The most reliable way to protect Microsoft SQL data is to schedule backups of the SQL Server to run at regular intervals. Scheduled backups ensure all SQL Server data is automatically saved, including the NetWorker server's client indexes and bootstrap file. The client indexes and bootstrap file are vital for restoring data to the SQL Server in the event of a disaster.

Scheduling the backups is done by an administrator using the NetWorker Administrator program and with a working knowledge of the NetWorker software. Scheduled backups can be configured to run at any time and use backup levels full (level 0), differential (level 1), and incremental (level 9). More information on backup scheduling and backup levels is available in the *NetWorker Module for Microsoft SQL Server Administrator's Guide*.

---

Another method to schedule backups is to use the NetWorker Configuration Wizard. The wizard is available as an add-on tool to easily create one or more client or device resources for basic scheduled backups of NetWorker client computers and Microsoft SQL Server data. The wizard integrates with this NetWorker Module by prompting users for information, such as the following:

- Backup type
- Objects to backup
- Recovery level
- Schedule preferences
- Microsoft SQL administrator account information

The NetWorker Configuration Wizard is installed separately from the NetWorker Module for SQL Server. When you choose to install the NetWorker Module for Microsoft SQL Server component of the NetWorker Configuration Wizard, this component allows you to use the NetWorker Configuration Wizard to configure scheduled backups of Exchange server data as well.

## **NetWorker Module for DB2 Universal Database (NMDB2)**

NetWorker Module software interacts with the DB2 Universal Database (UDB) Backup and Restore Utility to back up and restore database data.

### **NetWorker Module backup**

The NetWorker Module connects the DB2 UDB backup utility to the NetWorker server through the NetWorker libnsrdb2 library, which implements an X/Open Backup Services Application programming interface (XBSA API).

When a backup request is initiated, the DB2 backup utility interacts with NetWorker software through the XBSA API to coordinate a backup of the database.

When **nsrd** triggers a scheduled backup on the NetWorker server, **savegrp** executes the **nsrdb2** script instead of performing the standard save. The **nsrdb2** script prepares the environment and runs the **nsrdb2sv** binary. The **nsrdb2sv** binary invokes the DB2 backup utility to perform a database backup.

### **NetWorker Module restore**

When a DB2 UDB manager restore request is initiated, the NetWorker Module XBSA library translates the object names requested by the database into a format the NetWorker software understands and forwards it to the **nsrd** service on the NetWorker server. The media service, **nsrmmmd**, searches the NetWorker server online media database for the media containing the objects requested and restores the data to the database manager.

### **The db2uext2 program and log files**

The **db2uext2** program provided with the NetWorker Module performs transactional log backups and restores. The DB2 Universal Database must be configured for roll-forward recovery by enabling the USEREXIT database configuration parameter.

The database manager:

- Calls **db2uext2** to back up online logs to disk or NetWorker server.
- Requests **db2uext2** to restore logs during a database roll-forward.

You can also use the DB2 **archive log** command to call **db2uext2** to archive logs on demand.

---

**Important:** Only the **db2uext2** program supplied with the NetWorker Module is supported.

---

---

## **NetWorker Module for Informix (NMINF)**

NetWorker Module for Informix is an add-on module for NetWorker software that provides:

- Automated backup media management.
- Scheduling for ON-Bar, the Informix backup and restore utility.

NetWorker Module for Informix provides:

- The ability to integrate both database and system file backups, thereby relieving the burden of backup from the database administrator while allowing the administrator to retain control of the restore process.
- True “lights out” database storage management through automated scheduling, auto-changer support, electronic tap labeling and tracking.
- Support for local or distributed backup to a centralized backup server.
- High performance through support for multiple, concurrent high-speed devices such as DLT drives.
- Cluster support for high availability.
- Support for 64-bit Informix servers.

NetWorker Module for Informix, together with NetWorker software, provides reliable, high-performance data protection for local or distributed Informix Dynamic Server (IDS) databases. NetWorker Module for Informix integrates backup and restore procedures for Informix databases with the network-wide data protection solutions that NetWorker software provides.

NetWorker software, in combination with the NetWorker Module for Informix, provides a storage management solution that addresses the need for cross-platform support of SAP running on Windows NT, UNIX, or Linux operating systems.

### **NetWorker Module for Informix backup**

When a scheduled Informix backup is triggered by **nsrd** on the NetWorker server, **savegrp** executes **nsrdbmi** on the client instead of a standard save. The **nsrdbmi** command backs up data that is passed from ON-Bar and sends it to the NetWorker server through the NetWorker Module API. The results of the **savegrp** execution are sent to the NetWorker server and included in the savegroup completion report.

After a scheduled backup, NetWorker sends a record of the server’s bootstrap file to the default printer. This is a printed record of the dates, locations, and save set ID numbers for the server’s online indexes, required for restoring data. Keep the bootstrap printout on file as a quick reference in the event of a disaster, such as a disk crash or server malfunction.

Figure 8 on page 16 shows how data moves from the database server to the NetWorker server during an ON-Bar backup session.

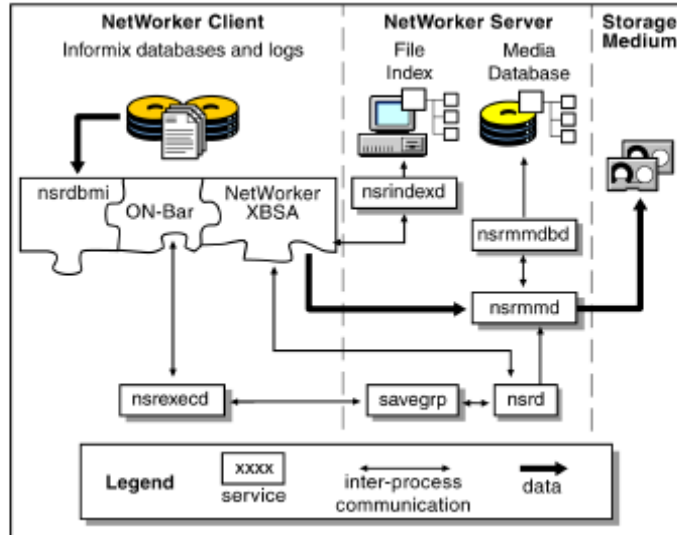


Figure 8. Backup initiated by NetWorker Module for Informix

### NetWorker Module for Informix recovery

When an ON-Bar restore request is initiated, the NetWorker Module API translates the object names requested by ON-Bar into a format understood by NetWorker and forwards them to the NetWorker server's **nsrd** service. The media service, nsrmmmd, searches the NetWorker server's online media database for the media containing the object(s) requested and recovers the data to the database server.

Figure 9 on page 16 shows how data moves from the NetWorker server to your database server during an ON-Bar restore session.

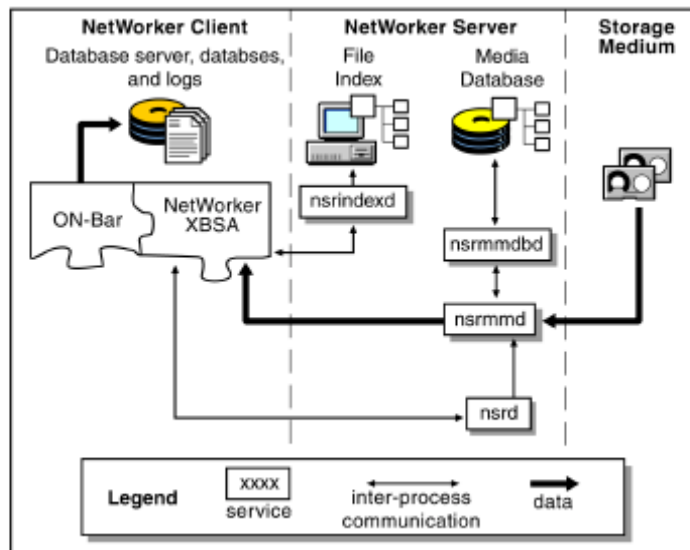


Figure 9. Data movement during a restore initiated by ON-Bar

---

## Split-mirror backup with SAP databases and NetWorker

SAP advises split-mirror backup when availability is an important consideration for the production system. For more information on classic split-mirror backup, refer to SAP Notes for your database of choice at the SAP Service Marketplace (<http://service.sap.com/notes>).

### *The split-mirror concept*

Classic split-mirror backup is typically done by creating a full replica of the database and importing it to another host, referred to as the backup mount host. The backup mount host will typically have the database software used for SAP installed. Once the replica is available on the backup host, it can be backed up.

### Method of backup with split mirror

The backup can be managed from the backup host as a database backup, using the appropriate NetWorker module for your database. An alternative is to merely treat the data as files to be backed up using NetWorker in a “classic” sense without the NetWorker database module.

### Options for creating the split-mirror replica

There are three alternative methods to accomplish the split-mirror copy for backup. The best technical solution will be determined by requirements and characteristics of any individual SAP enterprise environment and infrastructure. For more information tailored to best suit your SAP environment, contact your local EMC Sales Representative.

### **PowerSnap and NetWorker**

PowerSnap integrates with the NetWorker Module to provide split-mirror backup functionality integrated with your backup. Using PowerSnap with the NetWorker Module delivers the advantages and benefits of split mirror in a complete NetWorker solution.

A PowerSnap for SAP with Oracle backup is started by automatically or manually invoking the scheduled NetWorker backup group on the NetWorker server. **BRBACKUP** will start NetWorker Module’s **backint** program to integrate with **BRBACKUP** for Oracle. The **backint** program manages PowerSnap to perform the replica for backup. The end result is recorded in SAP’s tables by the BRTools and can be viewed from CCMS.

### **EMC Replication Manager and NetWorker**

Allowing Replication Manger (RM) to handle the replication technology means there is no need to maintain storage device groups for replicas. RM also manages storage resources for replication whether they are full-copy mirrors, snapshots, clones, or BCVs. RM uses a storage pool predefined for backup use. Replication Manager applies algorithms to select the best possible replica device pairings each time a replica is created.

Replication Manager can also be integrated with **BRBACKUP** to produce a replica of the database for backup. In this method, **BRBACKUP** can leverage existing or new RM software to provide the replica on the backup mount host and call NetWorker to stream the data to the backup device. **BRBACKUP** drives the process, and the end result of the backup is recorded in SAP tables by the BRTools. It is viewable using SAP CCMS.

### **Manually maintained scripting with NetWorker**

Creating a split-mirror replica can be done by means of a script or program that the manufacturer of the storage infrastructure provides to split and later resynchronize volumes.

---

## Split mirror with BRBACKUP for Oracle

With a split-mirror configuration, the volumes of the production database are paired with volumes for creating the replicas. **BRBACKUP** in this scenario runs on the backup mount host, where the backup is actually performed after the “third mirror” disks have been split and mounted. This is accomplished by adjusting the `init<SID>.sap` configuration file to include the following:

- `backup_type= online split`
- `split_cmd=full path and script file name to execute to create and mount replica`

In a typical scripting solution with split mirrors, **BRBACKUP** can be used to control the splitting of the volumes. **BRBACKUP** will drive the backup process and the success or failure of the backup will be recorded within SAP; viewable through SAP CCMS.

For more specific information about **BRBACKUP** and split-mirror backup, refer to SAP Notes 170607 and 378818 at the [SAP Service Marketplace](http://service.sap.com/notes) (<http://service.sap.com/notes>).

## New options for backup and recovery

SAP customers are enhancing their tape backup strategies with new disk-based options. By tapping the capabilities of disk, such as concurrent read/write and random access, SAP customers can complement their tape backup strategies to achieve faster and more reliable backup and recovery while continuing to use tape for what it does best—long-term storage that is easily transportable.

The storage industry and storage strategists have developed new approaches to data backup and recovery. New technologies, such as low-cost ATA disk and enhanced backup software, give IT, Basis, and database administrators new options for meeting heightened backup and recovery requirements.

The new backup and recovery options revolve around the emergence of low-cost disks. Taking advantage of greater capacity and lower pricing, storage vendors and strategists have been designing new backup solutions, such as backup to disk and tape library emulation, also known as virtual tape.

These new options promise a number of benefits over conventional tape backup:

- Faster data recovery, which gets the organization up and running sooner
- Quicker backup, which addresses the challenge of diminishing backup windows
- Increased confidence in the backup due to redundancy, RAID, and other high-availability features
- Greater reliability due to the better reliability of RAID-protected disks, as compared to tape

In short, the new disk-based backup options finally solve the problems that have been plaguing conventional tape-based backup for decades.

**75% of storage management is backup and recovery, yet 30% of all data recovery instances fail due to botched backups.**

— Anders Lofgren, Forrester Research, 2003

## Considerations when selecting backup-to-disk options

Backup to disk is an implementation that offers service-level improvements over the traditional tape backup process. With changing economics of disk technology, backup-to-disk solutions are affordable, and SAP customers are implementing backup-to-disk solutions as improvements to their existing tape solutions.

Backup to disk means writing to a disk file in a file system. The premise is all primary backups are targeted to disk. This enables significantly faster backups, and even more importantly, provides near-instantaneous restore when necessary.

This is possible because users can back up data at the write speeds of the target storage systems (CLARiiON or NS700). Disk drives are random-access devices and can immediately begin writing data.

---

In contrast, tape technologies require positioning and respond to a minimal data stream by excessive positioning, also known as *shoe-shining*.

## Advantages of backup to disk

Disk-based backup solutions offer fast and reliable backup and restore operations. Such solutions are complementary to physical tape libraries to provide new flexibility in designing backup and recovery scenarios for SAP.

Bringing disk into the tape paradigm adds definitive value; especially in the area of restore. Traditionally, one had to choose between fast backup or fast restore. In order to maximize backup throughput, avoid tape delays, and the wear and tear of tape drive “shoe-shining”; an environment requires tuning to keep the tape drives streaming during backup. Tuning involves a process of configuring the backup application to accept multiple streams of data, from multiple sources, and interleaves them into a datastream that is sent to a device (*multiplexing*). With multiplexing, the drives are kept spinning to provide the best performance with the least amount of impact on the tape drive(s). The price is paid when it is necessary to recover from this type of backup. To get the specific data needed from the backup, the application must “demultiplex” or unwind the data stream that is intermixed with multiple sources. While this is handled seamlessly by NetWorker, the recovery experiences a degraded performance. Recovery performance in the case of demultiplexing is directly proportional to how high of a degree the backup environment is tuned to provide fast backups.

In a backup-to-disk environment, data is recovered at the processing speed of the (CLARiiON or NS700) storage systems and is read in a contiguous manner, enabling users to gain benefits of protecting clients in parallel without sacrificing recovery performance. It no longer interleaves the datastreams to provide performance, eliminating the need to demultiplex data for a recovery. Users now can tune for both backup and restore performance. Also, tape mounting and unmounting delays are eliminated.

Whether writing to a CLARiiON LUN or an NS700 file system, multistreaming improves aggregate performance by allowing read and write of multiple streams of data to the same LUN or file system simultaneously.

Media-specific errors, including faulty tape media and tape mounting failures, which can cause backups to fail, are minimized. RAID protection of disk systems also ensures data availability in the event of a disk drive failure. CLARiiON and NS700 are highly available storage systems with no single point of failure. A result of the availability improvements is fewer full backups may be required, which saves network and CPU processing cycles, and improves total cost of ownership (TCO).

Tape handling is reduced, and in some cases eliminated. Maintaining sets of tapes from tape libraries can be problematic and require properly trained personnel. Disk does not require the tape handling/positioning.

## Disk-based backup: Two enterprise options

There are two enterprise options for disk-based backup: tape library emulation and backup to disk. The two approaches are very similar and use the same underlying disk and backup software. However, each has different characteristics that make it more appealing based on an organization’s specific requirements and situation.

One important distinction to consider with disk-based backup is the focus on operational backup and restore. Tape still delivers significant advantages in terms of the removability and mobility of the tape medium for purposes of offsite storage and ability for long-term retention.

Disk-based backups are intended to augment tape backup; allowing faster, more reliable, and more efficient backup and recovery of the content residing onsite. Disk-based backups can be cloned or staged to tape for shipment to an offsite location or for disaster recovery without impacting daily operations.

### Tape library emulation

Tape library emulation is a physical disk solution designed to emulate a tape library. The disk capacity appears to the backup software and applications as a physical tape library. From a backup point of view, the emulated tape library is seen as a standard backup target. The result is that no changes are required to

---

the backup software, the applications, or the data backup process. The only noticeable change is data backup and recovery happens significantly faster.

The primary advantages of emulated tape over physical tape are drastically improved backup and restore performance, significantly enhanced reliability, and improved manageability. Additionally, tape emulation can be substituted for physical tape library in a plug-and-play fashion. Emulated devices do not suffer from mechanical delays such as robotic tape picking, loading, and tape fast-forward or rewind. Searching for data on tape becomes an immediate operation through the advantages of disk's random-access capabilities.

Another advantage of tape emulation is the ability to create more tape drives than were previously available in the physical environment. This permits the backup administrator to redeploy backup jobs that may have been waiting for tape drive resources to become available. This can also eliminate tape sharing, which ultimately will optimize restore performance. More backup jobs can be run in parallel and complete in less time.

### **Backup to disk (B2D)**

Backup to disk, where disk is presented as a file-system-based volume, is similar to tape library emulation in many respects. It uses increasingly low-cost disk capacity to back up data more quickly and reliably. More importantly, restore of data is quick and reliable. In this case, the backup application treats disk as disk. This allows the enterprise to take full advantage of the inherent capabilities of disk, especially the capability of random access; most significant is the ability to read and write data on different parts of the disk at the same time.

Backup to disk generally requires some changes to the backup process. The backup software must be configured to see the disk as a backup target and treat the disk as disk to take full advantage of concurrent read/write and other disk capabilities. These changes are modest, and NetWorker supports a wide range of options.

Cloning, the process of making a mirror copy of stored data, can begin on any completed save-set even when the full backup process is still running. This is a unique feature of NetWorker Disk Backup Option (DBO). Traditionally, cloning could only begin after completing all save sets in a group. By taking advantage of concurrent read/write capability, backups and clones can finish in nearly the same time it normally takes to complete only a backup. Once cloned, if a restore from tape is ever necessary, data can be recovered directly to the affected clients in a single step.

## **Choosing a disk-based backup option**

To determine which approach is right for a given enterprise, decision-makers must look at the environment, operational concerns, and budget considerations. The advantages of disk-based backup are very compelling; it is almost impossible to make a wrong choice when adopting a disk-based approach to backup. Though either solution can likely fit into any environment and produce desirable results, organizations may find they favor one over the other based on some specific benefits.

Backup to disk (DBO) works best where organizations:

- Are comfortable and experienced in disk management.
- Want to take advantage of the flexibility and capabilities to maintain backup and production data on the same storage array, and exploit simultaneous read/write on the same file or virtual device.
- Have existing disk within an environment that can be allocated for backup (or have a storage array to which ATA drives can be added).
- Backup administrators need to restore recent backup with a high frequency level.
- Users want to implement the best performing backup/cloning operations.

Tape library emulation with CLARiiON Disk Library (CDL) approach works best where organizations:

- Are comfortable and experienced with the tape paradigm.
- Want to gain disk speed advantages but choose not to alter their standard tape management practices.

- 
- Have backup teams that cannot make traditional disk decisions, but do have a budget for backup media.
  - Want to extend the advantages of disk backup to their NAS NDMP backup infrastructure.

As noted earlier, disk-based backup is not intended to completely eliminate tape. Tape may still be required for offsite storage and for long-term archival purposes. For more information and assistance with selecting disk backup options tailored to best suit your SAP environment, please contact your local EMC Sales Representative.

## ***EMC recommendations for B2D***

Many variables can be set within EMC storage systems and NetWorker Modules to optimize backup-to-disk performance to meet the requirements of an SAP landscape.

EMC recommends using ATA technology to optimize backup performance with NetWorker.

Major advantages of backup-to-disk include the following:

- Faster backup performance
- Faster data recovery
- Improved media reliability and data availability
- Improved IT management and efficiency

To ensure optimum backup-to-disk performance with CLARiiON storage systems:

- Enable CLARiiON write cache.
- Use only RAID 3 striping for backup to disk with CLARiiON. RAID 3 has been optimized for read/write performance with cache enabled and ATA technology.
- Configure all RAID groups as 4+1 or 8+1 stripes.
- Ensure source file defragmentation.
- Limit the number of backup streams/RAID groups to no more than four or five.
- Set file system block and element sizes to their default settings.
- Disable software compression where possible.

## **A two-tiered backup approach (D2D2T)**

A complete backup solution should include data protection in a tiered backup approach by first moving data to fast disk, and then transferring to tape at a later time for offsite storage. This approach minimizes the SLA backup and recovery window in an event of a local system outage. Further offloading backup disk to tape provides preservation of data for offsite and legal compliance. This approach leverages the existing tape backup investment and local backup disks can be rotated and reused for backing up new data.

NetWorker software provides two ways to move backup to tape: cloning and staging. Cloning will copy the backup image to tape, leaving the original on disk. This creates an additional copy of the backup image for mobility. Staging, on the other hand, copies the backup image to tape and removes the original from disk to conserve space.

---

The staging policy is defined at the NetWorker server via the staging resource. The staging resource names and defines the NetWorker policy for moving backup data from one storage medium to another, and removing the data from the original location. Backup images are staged based on any or all of the following user-defined policies:

- Age
- Storage medium capacity limits (high and low watermarks)
- Backup size

## Conclusion

NetWorker simplifies and centralizes backup and recovery operations across a heterogeneous SAP environment. There are options for further optimization through use of backup to disk with cloning or staging to tape.

There are now choices of traditional disk, and virtual disk libraries to gain backup/restore advantages. This provides the ability to design backup and recovery scenarios that work best for any particular requirement or user skill set.

EMC NetWorker is a market-leading, globally proven solution for protecting business-critical information and provides a solid foundation for reliable enterprise storage management.

## References

- *NetWorker Module for DB2 Universal Database Administrator's Guide*
- *NetWorker Module for DB2 Universal Database Installation Guide*
- *NetWorker Module for Informix Administrator's Guide*
- *NetWorker Module for Informix Installation Guide*
- *NetWorker Module for Microsoft SQL Server Administrator's Guide*
- *NetWorker Module for Microsoft SQL Server Installation Guide*
- *NetWorker Module for SAP R/3 with Oracle Administrator's Guide*
- *NetWorker Module for SAP R/3 with Oracle Installation Guide*
- *NetWorker Administrator's Guide*
- *NetWorker 7.x DiskBackup Option*

For further information, see:

- <http://service.sap.com/ATG> (click Backup & Restore)
- [www.EMC.com/sap](http://www.EMC.com/sap)
- <http://help.sap.com> (search for split-mirror backup)