



# PROTECTING IT WITH NETWORKER

Aaron Kleinsmith  
Principal Technical Instructor  
EMC  
[aaron.kleinsmith@emc.com](mailto:aaron.kleinsmith@emc.com)

EMC<sup>2</sup>

## Table of Contents

Protecting the Modern Information Technology Environment .....	5
NetWorker Modules.....	7
Diagram 1: NetWorker Module Backups.....	7
NetWorker Module for Microsoft .....	8
NetWorker Module for Databases and Applications.....	8
Other NetWorker Modules.....	8
Configuring the NetWorker Modules.....	9
NetWorker Snapshot Management and Block Based Backup .....	9
When would these different snapshot technologies be used?.....	9
NetWorker Snapshot Management for NAS .....	10
NetWorker and Windows Bare Metal Recovery.....	10
Data Domain .....	12
Data Domain CIFS and NFS .....	12
Data Domain VTL.....	13
NDMP with NetWorker and Data Domain.....	13
Local NDMP Backup .....	13
Three-way NDMP Backup .....	14
Data Domain in a Three-way NDMP backup .....	14
Data Domain as the tape library to the NDMP Tape Server.....	14
2015 EMC Proven Professional Knowledge Sharing .....	2

Backup to non-NDMP device with NetWorker NDMP-DSA.....	14
Data Domain Boost .....	15
Data Domain Distributed Segment Processing and Client Direct.....	15
Diagram 2: Distributed Segment Processing with Client Direct .....	16
Clone Controlled Replication .....	17
Diagram 3: Clone Controlled Replication with offsite Storage Node.....	18
Interface Groups.....	18
DDBoost Database Modules .....	18
Avamar.....	19
Table 1: Supported file systems and applications by Avamar. ....	20
Avamar and NetWorker Integration .....	20
Avamar and Data Domain Integration.....	21
When to integrate Avamar? .....	21
Avamar Virtual Edition .....	22
NetWorker VMware Protection .....	22
Diagram 4: NetWorker VMware Protection .....	23
NetWorker VADP (Legacy) VMware Backup.....	24
NetWorker Guest-Based backup of VMware .....	24
Data Protection Advisor.....	25
Backup and Recovery Manager .....	26
2015 EMC Proven Professional Knowledge Sharing	3

Example 1: of the Backup and Recovery User Interface .....	27
Example 2: Backup and Recovery Manager Mobile Interface .....	28
Example 3: Backup and Recovery Manager Mobile Interface .....	28
EMC SourceOne .....	29
Backing up SourceOne with NetWorker .....	29
Table 2: EMC SourceOne backup scripts .....	31
Conclusion .....	33
Appendix .....	34

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

## **Protecting the Modern Information Technology Environment**

Every information technology environment is becoming more complex for backup and recovery and the features that keep the systems always alive and working mean that there is less idle time to copy the data. Backup strategies and software features are ever updating and changing to keep up and work towards an ideal situation where the systems we want to back up are no longer impacted during a backup window.

NetWorker<sup>®</sup> software was originally developed when backing up across a Local Area Network (LAN) was a new idea. At the time it was great to be able to back up data to a central place. When it was necessary to back up an application it was normally shut down and the inactive files were backed up through the file system back up. This would ensure any activity in memory was flushed to the disk and data wasn't changing during the back up.

Of course, over time this was no longer going to meet the business demand since access to the applications is often required 7 days a week, 24 hours a day. Many database vendors created methods for accessing the data in a consistent state while the applications were still up and running. NetWorker Modules were added to the base NetWorker file system backup framework and the Modules were designed to provide a specific tie-in to the database backup utilities and APIs provided by the database vendors.

Even this was not ideal since the host that was running the application is still part of the backup process using CPU, RAM, and other server resources to get the data copied off the systems and to central storage managed by NetWorker. Server-less backup, or a backup that does not even use application server resources would be ideal. To provide the best possible availability, we have all kinds of ways of replicating, mirroring, and capturing data from the applications while they are active. Backing up the mirrored copies instead of the primary would allow a server-less type of backup and some of the functions of NetWorker will allow this with other integrated operating system features, application features, disk or storage array features, virtualization options, or features in NetWorker.

There have never been more options, features, and supporting products from EMC that are part of or work with NetWorker. I find many IT administrators struggle to figure out how all the different pieces of EMC's Information Infrastructure products work together, and in this case, with NetWorker. One common scenario is that the administrators who will configure and manage NetWorker are not usually involved in the decision for purchasing NetWorker and don't get involved until it is time to manage the day-to-day backups. Additionally, NetWorker's licensing model has changed in recent years allowing existing customers to access additional features they weren't previously using or aware of.

The target audience for this article includes NetWorker Administrators who want to gain a better understanding of how NetWorker can address their client's needs. Clients such as database administrators, network administrators, server administrators and email administrators that manage the organization's enterprise environment, but not usually the backup infrastructure.

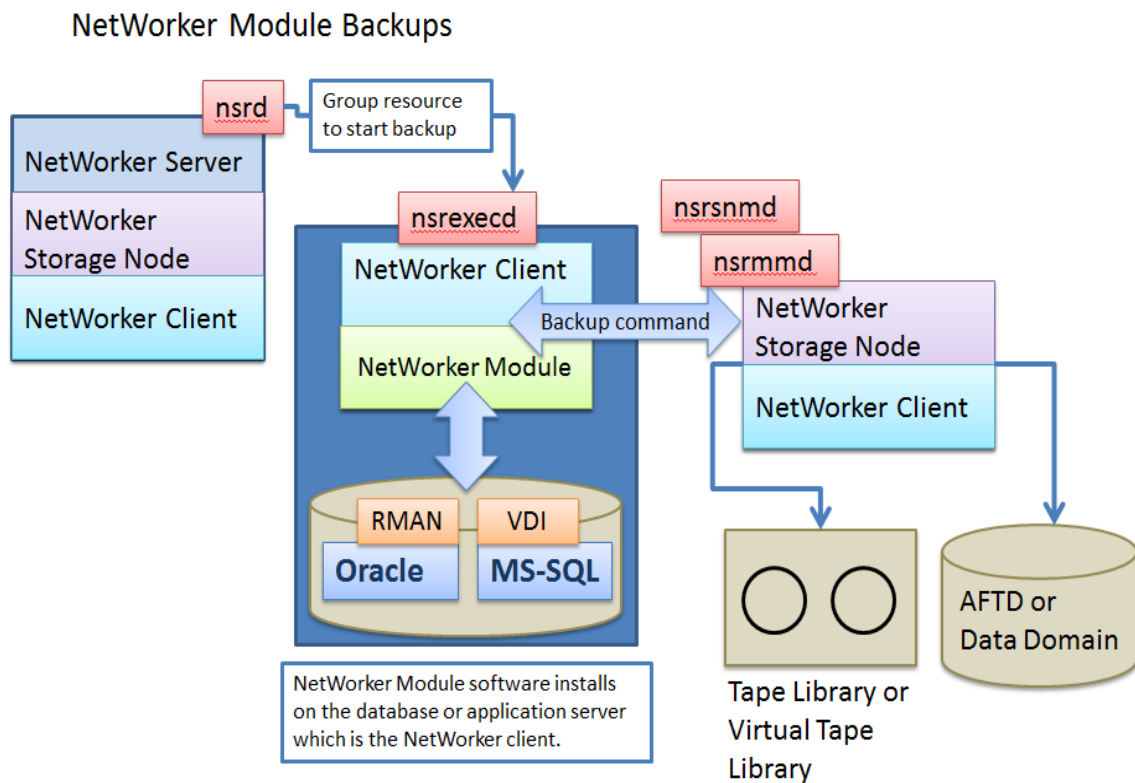
This article will help explain backup and recovery with NetWorker as an overview with many topics in the one place that could be consumed in a reasonable amount of time. It will serve as a starting point before consuming the more detailed guides available from the EMC support web site.

This article describes how NetWorker is used to coordinate backup and recovery activity for an enterprise environment using complimentary EMC products and add-on features such as NetWorker Modules, Snapshot Management, Block-Based Backup, Data Domain®, Avamar®, Data Protection Advisor, Backup and Recovery Manager, and EMC Backup and Recovery Appliance. It is not inclusive of all possible combinations that can be configured but covers some of the most common options. This article does not include detailed implementation, configuration, and management information that would already be included in other documentation or training available from EMC.

This article addresses what many of these features are and discusses why each fits for different solutions.

## NetWorker Modules

To protect backing up applications such as databases it is important that all data is captured in a consistent state at the time of backup so that when a recovery is performed, all files and transactions for a particular database are in-sync to a point of time of the backup or a point in time needed for recovery. Each database or application vendor has a different way of allowing the backup applications to copy the data. To meet all of the needs, NetWorker Modules are installed after the regular client installation and are called for backup instead of the usual 'save' command that would be used to backup file system data.



**Diagram 1: NetWorker Module Backups**

## **NetWorker Module for Microsoft**

The NetWorker Module for Microsoft Applications (NMM) is used to back up applications such as Microsoft SQL, Microsoft Exchange, Sharepoint, Hyper-V, and additional Active Directory backup and recovery support.

NMM originally were separate modules for each application such as NetWorker Module for SQL, NetWorker Module for Exchange, etc. When Microsoft introduced Volume Shadow Copy Services (VSS) there was a combined NetWorker Module for Microsoft Applications that utilized the shadow copy functionality of VSS to perform what NetWorker calls a snapshot of the application. Not all of the features from the older modules could be incorporated with VSS so eventually with the later versions (2.4, 3.0, 3.1, and now 8.2), the older functionality, VSS, and newer features have been put together in one Module. Now administrators can choose which systems they want to protect with VSS and which they would like to protect with the other APIs. The later 8.2 version was the next version after 3.1 in order to 'jump' the version number to match the NetWorker server, storage node, and client version. For example, there were no version 4, 5, 6, 7 released.

## **NetWorker Module for Databases and Applications**

The NetWorker Module for Databases and Applications (NMDA) is used to back up applications such as Oracle, DB2, Informix, Lotus Domino, Sybase, and MySQL. The latest version of NMDA is also version 8.2 which is named and released to match the latest NetWorker version.

## **Other NetWorker Modules**

The NetWorker Module for SAP has support for backing up SAP R/3 and SAP HANA. The NetWorker Module for MEDITECH has support for backing up a MEDITECH environment.

Please check the NetWorker Software Compatibility Guide for all supported applications and operating systems and different NetWorker versions that work together.



## **Configuring the NetWorker Modules**

To fully support the applications when configuring NetWorker Modules requires configuring a client to back up the file system and operating system files (non-database, non-application data) that will not be captured by the NetWorker Module backups. With the configuration of this client resource, a Directive can be configured and assigned to the client resource in order to exclude the database or application data that will be backed up with the NetWorker Module backup. At least one other client resource is configured to back up the database or application and it is used to call a special NetWorker Module command configured in the Backup Command attribute in the client resource. This command—the save set attribute and Application Information attributes in the client resource—can be configured with or without the Client Configuration Wizard to perform a proper backup. Please check the associated Administration Guide for the NetWorker Module being used as each database or application does have unique features and requirements for these settings.

## **NetWorker Snapshot Management and Block-Based Backup**

Two current features of NetWorker 8.1 that were additional NetWorker Modules in previous versions are now included as part of the base NetWorker client software. These features are the NetWorker Snapshot Management (NSM) and Block-Based Backup (BBB). NSM was formerly the PowerSnap Module and BBB the SnapImage Module. In NetWorker 8.2, another new related feature is NetWorker Snapshot Management for NAS.

These snapshot technologies help capture point-in-time snap or copy of the backup data. NetWorker Module for Microsoft Applications (NMM) also have some snapshot capabilities for certain applications through VSS writers.

### **When would these different snapshot technologies be used?**

NSM is used for performing snapshots of Windows servers when the storage volumes are on VMAX<sup>®</sup> storage arrays, VNX<sup>®</sup> Block storage arrays, and RecoverPoint Appliances. If you are planning to use the NSM capability of NetWorker, please reference the Snapshot Management Integration Guide for details on capabilities and limitations to using the feature.

BBB can be used to perform snapshot and block-level full and incremental backups of Windows servers. This type of backup is particularly useful if backing up a Windows file server and the server has many files or small files. When the save command does a backup of the Windows file system, it reads through the file system one folder and file at a time. If there are a lot of these objects, it can take a long time to do the backup and capture the metadata to be recorded in NetWorker's client file index.

## **NetWorker Snapshot Management for NAS**

NSM for Network Attached Storage (NAS) is used to perform snapshots of EMC Isilon<sup>®</sup>, VNX, VNX2, VNXe, and NetApp storage systems.<sup>1</sup> When planning to use the NSM for NAS capability of NetWorker please reference the Snapshot Management for NAS Devices Integration Guide for details on capabilities and limitations of using the feature.

In NetWorker, the Snapshot Policy resource and Client Configuration Wizard can be used to help configure some of the snapshot features.

## **NetWorker and Windows Bare Metal Recovery**

NetWorker has support with the NetWorker client software to back up and perform a Bare Metal Recovery of Windows 2008, Windows 7, Windows 8, Windows 8.1, and Windows Server 2012.

The NetWorker Windows Bare Metal Restore (BMR) will allow a backup of a physical Windows system through Windows Automated System Recovery (ASR) and Microsoft Volume Shadow Copy Services (VSS) Writers. The NetWorker BMR backup supports file system backup and recovery to be able to recover the operating system to the point-in-time of the backup with an automated procedure. If the system also has an application support by the NetWorker Module for Microsoft Applications (NMM), it will still be required to back up and then restore the application data in the method supported by the application and module. This can be recovered after the operating system and file systems and restored back from the BMR backup.

NetWorker can perform a BMR of newer Windows x86 and x64 systems using the NetWorker client software. Please check the latest NetWorker Software Compatibility for the exact version or systems that are supported.

When performing a full backup of a supported NetWorker client, there is an additional save set that is backed up called DISASTER\_RECOVERY:\ which will include the information needed to rebuild Windows system back to its state at the time of the backup. The data in the DISASTER\_RECOVERY:\ save set, plus volume information for those file systems that contain operating system data (critical volumes) are also backed up. They are used for the Offline System State Recovery (OSSR) to rebuild the volumes as part of the BMR recovery.

Specifically, a full backup of the DISASTER\_RECOVERY:\ save set contains:

- All Critical Volumes
- WINDOWS ROLES AND FEATURES
- System Reserved partition
- UEFI partition (if available)

When an incremental backup of the DISASTER\_RECOVERY:\ save set takes place, the Critical volumes are backed up at an incremental level while the other save sets are backed up at level full.

When a recovery of a system is needed, Windows BMR ISO image will need to be downloaded from the support.emc.com web site. This image will contain the necessary files to boot and create a WinPE system with a NetWorker Bare Metal Recovery Wizard that is used and started as part of the restore. The NetWorker and Windows files that are part of the ISO image will allow the system to boot and run the NetWorker System Recovery Wizard for the administrator to input settings around the NIC setup, hostname, NetWorker server, restore critical disk, and choose which backup to restore.

Additional details for options available for the NetWorker BMR can be found in the NetWorker Administration Guide. The NetWorker Procedure Generator that is accessible through the Solve Desktop at support.emc.com can be used to generate specific steps for performing a recovery to the same or similar hardware or—if the backup was from a supported hardware platform—there is an option to perform a physical to virtual (P2V) recovery. The compatibility guide has a list of supported hardware and software options for P2V. This might be a good option for migrating a physical server to the virtual environment.

If the servers are virtualized through VMware or Hyper-V, an image-level backup and recovery might be another strategy to deploy to protect the operating system. VMware virtual machines can be backed up with the NetWorker VMware Protection feature. Hyper-V can be backed up with the NetWorker Module for Microsoft.

## **Data Domain**

Data Domain is an EMC storage appliance that is targeted to store backup and archive data. Available in various models and sizes to meet backup storage needs at different places within an organization, Data Domain deduplicates data that it stores and does a very good job of receiving data and deduplicating and compressing that data as it is received. This deduplication process is referred to as in-line since the data is deduplicated before it writes anything to disk which will reduce the amount of writing and reading of the disk. Performing as much work as possible with CPU and RAM enables the overall deduplication process to complete quickly and allow the new backup data to be replicated off-site to another Data Domain sooner.

Each Data Domain storage appliance is made up of a head and disk. The head has the Data Domain Operating System (DDOS) installed and the disk will hold the backup and archive data in a file system unique to Data Domain. This file system is a logging type file system, referred to as the Data Domain File System (DDFS). The head and disk together make up the system as a whole, typically referred to as a Data Domain Server (DDS) or Data Domain Restorer (DDR).

Any data to be written on the Data Domain is passed through the head via Ethernet or Fibre Channel before being written to the DDFS. Backup and archive data can be passed to the DDR through protocols such as CIFS, NFS, VTL, NDMP, and DDBoost. Data Domain can be purchased with Ethernet or Fibre Channel connections, or both, depending on the requirement.

## **Data Domain CIFS and NFS**

With CIFS (Windows) or NFS (Unix/Linux), a NetWorker Storage Node can use an Advanced File Type Device (AFTD) to transfer data to Data Domain. The AFTD is known by device type `adv_file` in the NetWorker Management Console (NMC) GUI and can be set up by using the Device Configuration Wizard found when right-clicking the

Devices folder in the Devices area of NMC. The options selected would be the same to configure other CIFS or NFS type devices in NetWorker.

### **Data Domain VTL**

The Virtual Tape Library (VTL) feature of Data Domain enables a NetWorker Storage Node to transfer backup data through Fibre Channel by interfacing to Data Domain through operating system drivers for LTO1, LTO2, LTO3, or LTO4 tape drive types. The VTL feature is ideal for providing deduplication and disk storage benefits to software or applications that have specific requirements for using tape libraries and tape drives. An environment that has been using and will continue to use tape drives on a SAN may also choose to use the VTL feature to ease the transition to new backup hardware such as Data Domain. Features, procedures, configuration, or scripts in the backup software—such as NetWorker—that have been designed around tape can continue to be used.

### **NDMP with NetWorker and Data Domain**

Network Data Management Protocol (NDMP) was designed to back up NAS devices where NetWorker client software could not be installed. The primary backup device type when NDMP was designed was tape. The VTL feature of Data Domain is also a good fit for backing up NDMP data.

NetWorker can back up NAS systems such as \Isilon, \VNX, \VNX2, \VNXe, NetApp, etc. to a tape library using the NDMP protocol. The EMC NetWorker Hardware Compatibility Guide contains a full list of supported NAS devices.

NDMP has its own terminology to describe the roles in a backup environment. The NAS device that is connected to the tape library is known as the NDMP Tape Server. The NAS device to be backed up is considered the NDMP Data Server. The NetWorker software is considered the Data Management Application (DMA).

### **Local NDMP Backup**

A backup configuration where the NAS device is both the NDMP Data Server and NDMP Tape Server is referred to as a Local Backup since the tape library would be accessible directly to the NDMP Data Server without sending data to another system connected to the tape library. When configuring within NetWorker, the NDMP Data Server is

configured as a NetWorker client and the NetWorker storage node – similar to a Dedicated Storage Node configuration.

### **Three-way NDMP Backup**

A three-way backup for NDMP is when the NDMP Data Server and NDMP Tape Server are different systems. It is referred to as a three-way backup configuration since the tape library would be connected to a different system than the NDMP Data Server (NetWorker Client) or the DMA (NetWorker Server). This third system in the three-way example is configured as a NetWorker Storage Node in the NetWorker Management Console GUI using a storage node type of 'ndmp' instead of scsi or silo-type storage node.

### **Data Domain in a Three-way NDMP backup**

A Data Domain Server can be configured to perform the role of a NDMP Tape Server in a three-way NDMP backup. This feature is configured through the 'ndmpd' command on Data Domain and adding virtual tape library devices to be used for the NDMP backup to the Tape Server access group. A VTL-created library can be made available to the Data Domain as an NDMP Tape Server and the backup data is passed from the Data Domain to the Virtual Tape Library internally. The backup data is sent to the Data Domain through the NDMP protocol from the NDMP Data Server (NetWorker client) through a Data Domain Ethernet connection.

### **Data Domain as the tape library to the NDMP Tape Server**

A second option for backing up NAS devices with NetWorker and Data Domain is to have Data Domain emulate a tape library through the VTL feature and present the tape library and tape drives across the Fibre Channel SAN to another NAS device that can be the NDMP Tape Server in the three-way backup. In this setup, the Data Domain is only discovered as a tape library attached to the NDMP Tape Server. NetWorker is not aware that the tape library is actually a Data Domain Server.

### **Backup to non-NDMP device with NetWorker NDMP-DSA**

A third option for performing NDMP backups with NetWorker is to use the Data Service Agent (NDMP-DSA) function of NetWorker. A NDMP-DSA backup with NetWorker can back up the NDMP data to a non-NDMP device. This type of backup is invoked by using

the “-M” (upper-case m) switch on the nsrndmp\_save command when configuring the NAS client resource. If using this special option on the nsrndmp\_save command the backup can be directed across the Local Area Network (LAN) to another non-NDMP storage node (Windows, Linux, Unix for example) and stored on a typical supported device type in NetWorker such as an Advanced File Type Device. More details on the different kinds of support for NDMP backups can be found in the NetWorker Administration Guide.

## **Data Domain Boost**

NetWorker can also write data to the Data Domain using the Data Domain Boost (DDBoost) protocol. The DDBoost protocol was designed to provide a method of transferring the backup data across Ethernet more efficiently than CIFS or NFS. Starting with DDOS version 5.4 and NetWorker 8.1, the DDBoost protocol can also be passed across a Fibre Channel connection. The DDBoost protocol is also supported by the Avamar backup software which is discussed in the next section

The DDBoost protocol is installed with the NetWorker client software and NetWorker Module software. A NetWorker Server, Storage Node, or Client can write to a Data Domain through the DD Boost protocol since they all contain the NetWorker client software components. Additional features are available with the DDBoost protocol that are not available with the other protocols described in the next sections.

## **Data Domain Distributed Segment Processing and Client Direct**

One of these features is Distributed Segment Processing (DSP) which allows Data Domain to distribute some initial deduplication work to the NetWorker clients and storage nodes. Having some of the work performed by these systems enables the protocol to determine which data is duplicate data before sending it across the network for backup. This is accomplished through the NetWorker client and NetWorker Module software by enabling a setting on the client called “Client Direct”. With the NetWorker Client Direct feature enabled and backing up to a Data Domain with DSP enabled, the client and Data Domain will work together to accomplish deduplication at the client level.

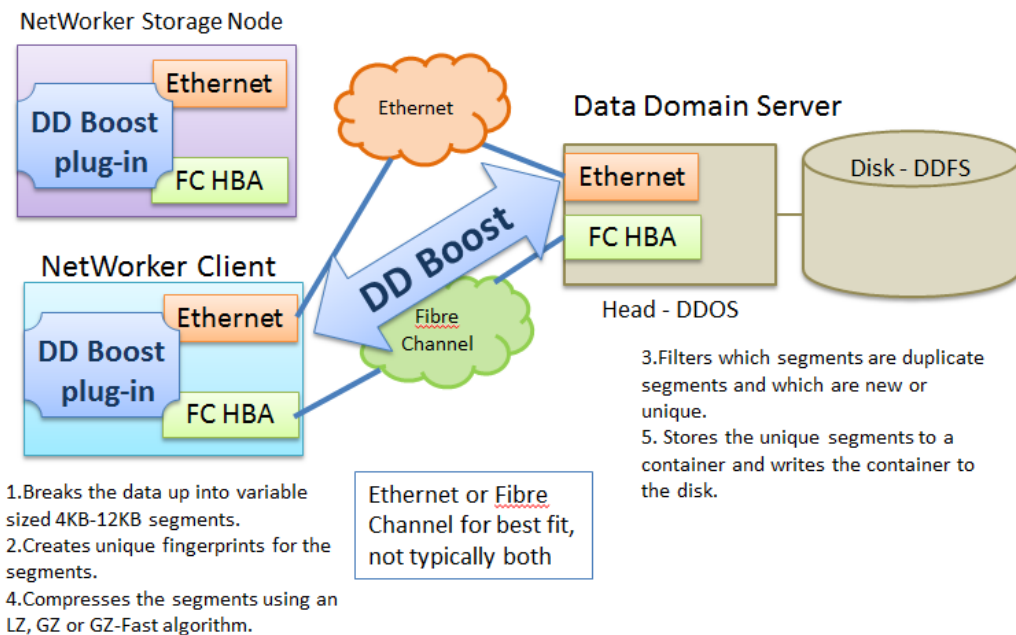
This is merely a simple and high-level description of what is happening during a backup to Data Domain to help explain how a DDBoost backup can help speed up the backup

window. The overall deduplication process that Data Domain performs has five main steps to storing the backup data that are considered part of the Stream-Informed Segment Layout (SISL) of Data Domain:

1. Breaks the data up into variable sized 4KB-12KB segments.
2. Creates unique fingerprints for the segments.
3. Filters which segments are duplicate segments and which are new or unique.
4. Compresses the segments using an LZ, GZ, or GZ-Fast algorithm.
5. Stores the unique segments to a container and writes the container to the disk.

Storage Node is still needed for nrsnmd and nsmmd processes to manage the DD device and for tracking information to be sent to NetWorker Server.

### Distributed Segment Processing



**Diagram 2: Distributed Segment Processing with Client Direct**

In a scenario where the DSP/Client Direct backup is being done, steps 1 and 2 are done on the client and then the fingerprints are passed back to the Data Domain to be sorted (step 3). The Data Domain reports back to the client the segments which are duplicate segments and which are new unique segments. The client will then compress only the new unique segment (step 4) and send those to the Data Domain saving bandwidth and



ultimately speeding up the backup window since the duplicate segments do not need to be sent. The Data Domain system will need to build the containers and write them to the Data Domain Filesystem (DDFS).

### **Clone Controlled Replication**

Cloned Controlled Replication (CCR) is another feature of DDBoost that is useful when integrating NetWorker and Data Domain through DDBoost. With CCR, replication of NetWorker save sets between two Data Domain systems can be started, monitored, and tracked through cloning operations in NetWorker. Each copy of the save set—the original and clone copies—will have their own identity in the NetWorker media database and can also have different retentions assigned.

To configure in NetWorker each of the Data Domain systems need to be registered in NetWorker and the source Data Domain labeled to a backup type pool with the destination Data Domain labeled into a backup clone type pool. Any of the cloning methods—such as the Clone resource, Clones option in the Group resource, nsrclone command, or manually cloning—can be used to start the process.

If using the Clone resource, a schedule can be configured of when to start the cloning. If choosing the Clones check box in the Group resource and choosing a clone mode to clone after save set completion as soon as a save set is backed up, NetWorker will begin the cloning and therefore replication between the two Data Domain systems. The data is deduplicated at the source Data Domain and passed directly from the source Data Domain to the destination where the storage node will not be in the data path of the operation. The replicated data will pass through Ethernet connections on the Data Domain systems.

## Clone Controlled Replication

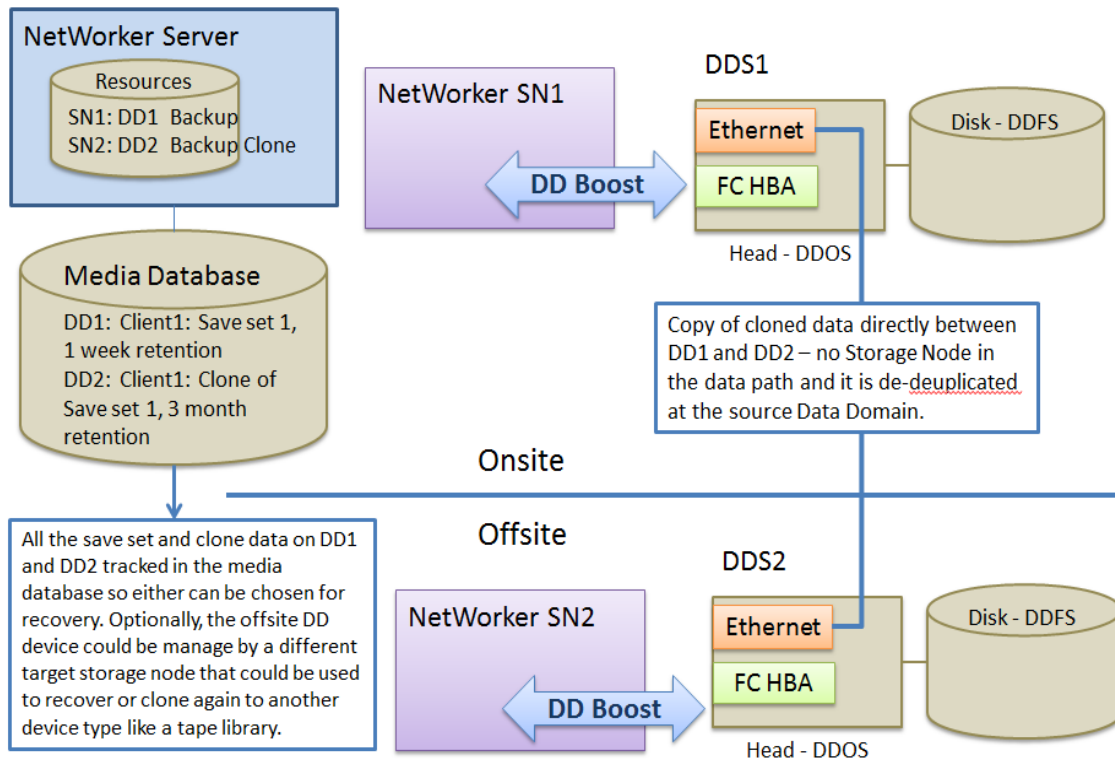


Diagram 3: Clone Controlled Replication with offsite Storage Node

### Interface Groups

Another benefit of DDBoost is Interface Groups that can be used to perform load balancing and failover with the network interfaces. Having DDBoost software code installed with the backup software enables the DDBoost code to handle the traffic with its own Interface capabilities without having to rely on the operating system or switch options.

### DDBoost Database Modules

There are also options for DDBoost to be used with other products including databases that can be backed up with NetWorker Modules. So why use NetWorker Modules with DDBoost backup? Whether an environment should use NetWorker or DDBoost modules alone to back up a database will depend on which administrators will be responsible for the backups. Will it be the NetWorker Administrators that need to meet the recovery point objective (RPO) or the database administrators in charge of meeting the RPO?

In some cases it might be easier to use a DDBoost plug-in installed directly on a database server such as Oracle, SAP, or Microsoft application. Be aware that those database backups will be managed separately from other backups in the organization and it may not be easy to monitor, troubleshoot, or manage those backups if they are configured outside of the NetWorker environment. Additionally, all the same tools we have for NetWorker may not be available. However, if using NetWorker there may be additional features with the NetWorker Modules and the monitoring and management of NetWorker in the integrated framework that we won't get from having the database interact directly with Data Domain. Also, NetWorker has an extensive list of supported device types such as tape libraries that can be used as another target for cloning backups to and from a Data Domain system.

## **Avamar**

Avamar is an EMC backup product that deduplicates the data that it stores. Avamar uses data deduplication to identify and store unique sub-file data objects only once. The global data deduplication that Avamar uses ensures that data objects are only backed up once across the entire Avamar backup environment.

Avamar uses disk to store the backup data and takes full advantage of inherent hard-disk characteristics. It can efficiently compress and reduce the amount of data sent over the network from a client and significantly reduce overall disk space the backups are using. Avamar supports a wide variety of client types and applications.

For the latest supported list, please check the EMC Avamar Compatibility and Interoperability Matrix. Here is a sample list of some of the supported file systems and applications supported.

File system backup	Application backup
Free BSD	Lotus Domino
HP-UX	Microsoft Exchange
IBM AIX	Microsoft Hyper-V
Linux	Microsoft Office SharePoint Server (MOSS)
Mac OS X	Microsoft SQL Server
Microsoft Windows	NDMP for NAS devices, including EMC Celerra IP storage systems and Network Appliance filers
Microsoft Windows Volume Shadow Copy Service (VSS)	Oracle
SCO Open Server	SAP with Oracle
SCO UnixWare	Sybase ASE
Oracle Solaris	
VMware	

**Table 1: Supported file systems and applications by Avamar.<sup>2</sup>**

An Avamar client is a system that has the client agent installed on it which enables that client to be backed up. An Avamar Server (or Avamar Grid) is made up of a Utility Node and multiple Data Storage Nodes. These are clustered Linux nodes that are used to manage the backups and receive the backup data. The Utility Node is the point of entry for managing the Avamar Server. The Avamar data server is made up of the Data Storage Nodes that each run a gsan process.

### **Avamar and NetWorker Integration**

Avamar can be used with NetWorker to provide deduplication functionality to NetWorker backups. NetWorker can be used to track the clients, schedule backups, and perform recoveries, while Avamar provides the deduplication technology to these backups. In a solution using NetWorker and Avamar together, the backup data is still stored within the Avamar disk storage in a deduplication node but the resources in NetWorker are used to configure the backup.

Avamar can also be used along-side NetWorker where it can schedule and manage its own backups and recoveries independently. Avamar does a very good job of identifying duplicate files and data chunks at the client level through information it keeps on the clients. It can identify files or data chunks before sending any information across the

network back to the central Avamar server. It uses a hash cache to store the unique meta-data to compare to when a backup occurs.

### **Avamar and Data Domain Integration**

The storage in an Avamar solution can be extended by sending some of the backup data to a Data Domain system which will allow the data to be further deduplicated on the back end. Depending on the model, a Data Domain system can have much more storage than an Avamar Grid alone. Therefore, a single Avamar Grid can end up storing some of its larger data content such as Oracle, Microsoft Exchange, Microsoft SQL, and VMware images over to Data Domain. Additionally, this data might deduplicate against other data on the Data Domain that has come from other kinds of backups such as NetWorker. Avamar communicates and sends the backup data to Data Domain through the DDBoost protocol which is installed automatically with the Avamar software.

### **When to integrate Avamar**

Avamar can be implemented as its own stand-alone back up product with many options for scheduling backups of many supported operating systems and applications. While there is a function to 'tape-out' some back up data, overall, the backups that Avamar collects will be written to the Data Storage Nodes (gsan) or Data Domain through DDBoost. Avamar does not have the same kind of support for tape library and virtual tape library backups as NetWorker. It does however, do an excellent job of backing up data across very slow network links and is used quite often to get a backup copy of data off a remote server or client system that would have previously been impossible. Avamar has continued to add supported data and application types.

If NetWorker is a well-established product in an environment and it is discovered that some of the client-side deduplication features of Avamar are needed, it could be integrated with NetWorker. In many cases Avamar and the systems that will require an Avamar backup can be configured for backup outside of NetWorker. Thus, if any issues occur with one of the products (NetWorker or Avamar), troubleshooting and maintenance such as upgrading can be handled individually.

## **Avamar Virtual Edition**

The Avamar Virtual Edition (AVE) allows the Avamar backup solution to be implemented within a VMware infrastructure and virtualizes the entire Avamar Grid within the virtual environment. This enables Avamar to be easily deployed in a virtual environment that has already been configured.

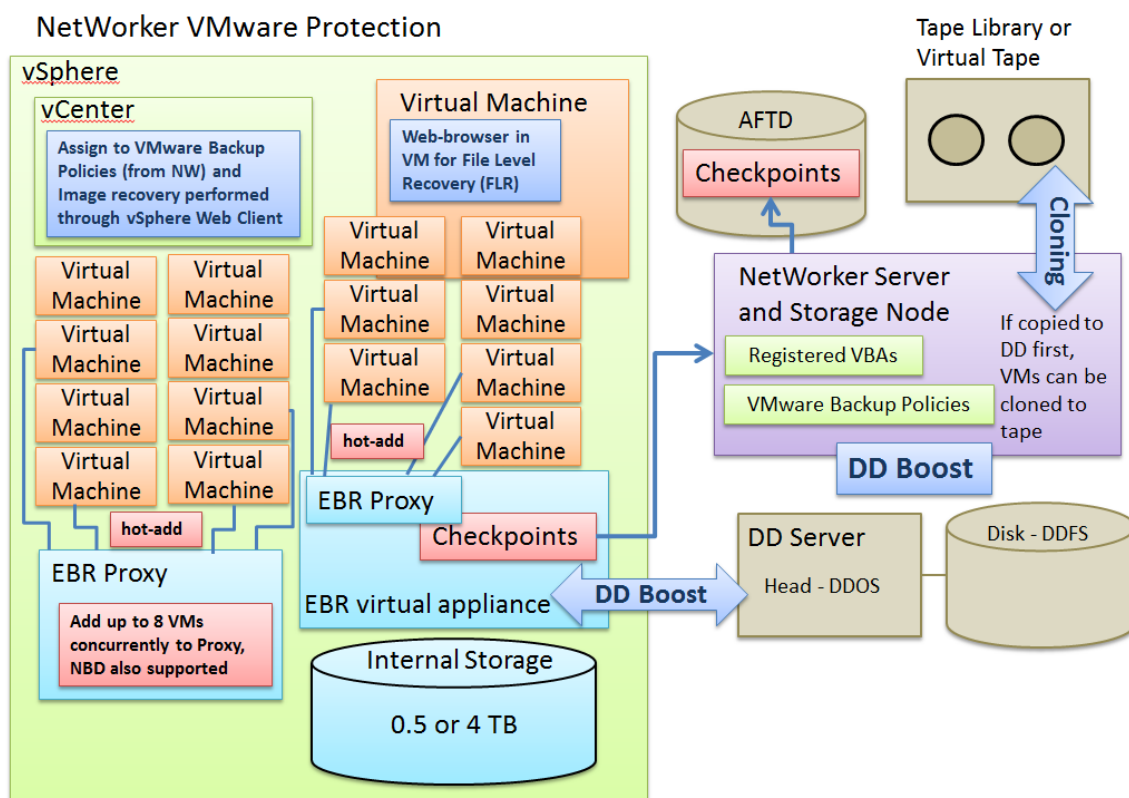
The AVE and NetWorker integration are used as the underlying technology for performing deduplication backups of VMware virtual machines with the NetWorker VMware Protection solution which is discussed in the next topic. In this newer option, for NetWorker to back up VMware, the function of being able to deduplicate with Avamar is leveraged in a seamless marriage of NetWorker and Avamar through VMware Backup Policies in NetWorker and a VMware Backup Appliance (VBA).

## **NetWorker VMware Protection**

A VMware Backup Appliance is a virtual appliance that can be installed into the vSphere environment and is used to back up and restore files, directories, and backup images using vStorage API for Data Protection. The VBA that EMC has for NetWorker is the EMC Backup and Recovery Appliance (EBR). Having a hardware-independent option for Avamar such as AVE has enabled seamless integration of Avamar backup and recovery into NetWorker VMware Protection through the VBA. There is a VBA available from EMC for NetWorker—the EMC Backup and Recovery Appliance (EBR). Starting with NetWorker 8.1 and 8.2, this is the preferred method to perform VMware virtual machine backups.

The EMC Backup and Recovery Appliance is an .OVA template that can be downloaded from [support.emc.com](http://support.emc.com) and deployed into the vSphere environment. The EBR virtual appliance contains scripts upon initial configuration that will seamlessly install and set up the needed Avamar and NetWorker components to be used by NetWorker to perform deduplicated backups to internal disk (gsan) or a Data Domain system through DDBoost. Proxy agent .OVA templates can also be deployed in the VMware environment to back up additional virtual hosts and increase the amount of concurrent backups that happen at once. The NetWorker VMware Integration Guide, available on [support.emc.com](http://support.emc.com), has additional details on how many concurrent backups and hosts can be deployed for the different supported versions of NetWorker.

EBR can be downloaded and deployed as a 0.5 TB image or a 4 TB image. The idea here is that backing up through DDBoost to Data Domain may not require as much internal storage; the 0.5 TB image may be all that is needed. On the other hand, If Data Domain is not an option for storing the backup data, 4 TB may be needed. A benefit of using Data Domain as a storage target for backups is that once NetWorker and the EBR have these save sets backed up to the Data Domain, NetWorker will be able to clone the data to other supported device types in NetWorker such as a Tape Library. If internal storage or Data Domain are not options for storing this backup data, there are still two other options that are supported.



**Diagram 4: NetWorker VMware Protection**

The EBR virtual backup appliance is installed into the vSphere environment and during initial setup, is registered with vCenter and NetWorker. In the Configuration area of the NetWorker Management Console there is a VMware Backup Appliance option in the navigation tree where details about the VBA can be viewed. There is also a resource type that can be created—VMware Backup Policies—which can include VMware actions

such as backup and cloning the required internal gsan or external Data Domain pool and an associated VBA. These backup policies are pushed out to the VBA in the vSphere environment where a VMware administrator can then assign virtual machines to the policies. There is no need to create client resources in NetWorker for backing up these virtual machines. Image recovery and file-level recovery for the virtual machines is managed from the EBR interface.

## **NetWorker VADP (Legacy) VMware Backup**

The older VADP (vStorage API for Data Protection) configuration for backing up virtual machines that was introduced in NetWorker version 7.6 and 8.0 is still supported as of now with NetWorker version 8.1 and 8.2. The configuration for this option within NetWorker is labeled with (Legacy). When configuring a backup, each of the virtual machines to back up need to have their own client resource in NetWorker. This is not a necessary step with the newer NetWorker VMware Protection feature. Each of the virtual clients will need to be associated to VMware Proxy which is also defined in NetWorker. The client resources for these systems can be defined through the NetWorker Client Configuration Wizard. If backing up using the legacy VADP method, the backups will be directed to standard supported NetWorker device types through the pool and device configuration. This type of backup would not require installing a NetWorker client agent in the virtual machine for performing a backup or an image-level restore. If needing to restore a select file or directory to a virtual machine, the NetWorker client would need to be installed into the destination virtual machine to which the file is being restored.

## **NetWorker Guest-Based backup of VMware**

Guest-based backup is a third option for backing up virtual machines. The guest-based backup has been available for quite some time since the premise is simple. The guest-based backup is configured by installing the NetWorker client software on the virtual machine and then creating a client resource in NetWorker. In this scenario, the virtual machine is treated similarly to a physical server that NetWorker would be backing up. It would not require any kind of specialized permission or setup within the vSphere environment. For this reason, it may be well-suited for those NetWorker Administrators that may not have access or knowledge of VMware. However, there are a lot of benefits to performing backups with the NetWorker Protection for VMware such as deduplication



and off-loading the backup load to a proxy host that would not happen with the guest-based backup.

One case where it is still recommended to perform a guest-based backup is when a NetWorker Module is required to back up an application that is on a virtual machine. A consistent point-in-time backup through the recommended database or application vendor method is important to ensure nothing is lost and the recovery is supported. Performing a virtual machine image backup will back up the image at a specific point-in-time. But when restoring, the applications on the image are restored back to a crash-consistent state which may not be sufficient to properly protect the application.

## **Data Protection Advisor**

Data Protection Advisor (DPA) is a collection and reporting software from EMC that can collect and report on a large number of applications and hardware components. The product started primarily targeting backup software, but includes much more than that now.

DPA software can help meet compliance and audit requirements, monitor and track backup failures, capacity planning, and trending to help meet agreed upon service level agreements (SLAs). Having the information that DPA provides will help organizations get the information they need for making decisions all in one place instead of having silos of information within each of the different products within an environment.

Some of the products that Data Protection Advisor can collect and report on include NetWorker, Data Domain, Avamar, NetBackup, BackupExec, Tivoli Storage Manager, Data Protector, CA ArcServe, CommVault, ACSLS, Oracle RMAN, Tape Libraries, VTLs, RecoverPoint, Vplex, SRDF<sup>®</sup>, TimeFinder<sup>®</sup>, MirrorView<sup>®</sup>, Oracle, Microsoft SQL, PostgreSQL, VSC, MSCS, Symmetrix<sup>®</sup>, VNX, NAS, IPSwitch, XSIGO, and FCSwitch.<sup>3</sup>

DPA software is comprised of three main components; the Application Server, the Datastore, and the Data Collection Agent. It uses the REST API to interface with different components and a PostgreSQL database to hold the database data it collects.

The DPA web console is a browser-based GUI which is used to monitor, analyze, and provide alerts and reports on backup and replication environments. There are dashboard

views to allow a single screen view of the status of the systems or software it is monitoring. To monitor NetWorker and generate reports related to NetWorker backups, a Data Collection Agent can be installed on the NetWorker server or collected remotely through a proxy server.

Once information is collected from NetWorker, DPA stores it in a PostgreSQL database that can be accessed to provide reports. Many reports are included and customized reports can also be created. The reports and alerts related to NetWorker enable administrators to stay informed and aware of any potential problems, all in one place where reports and alerts for other backup and replication products can be reported on. Another benefit of having access to the reports is the ability to easily generate information needed to answer questions by management and clients or to meet compliance requirements. Performance tuning is also aided by being able to see more performance-related information and a history of the performance of the backups and deduplication with more detail than can be obtained from NetWorker alone.

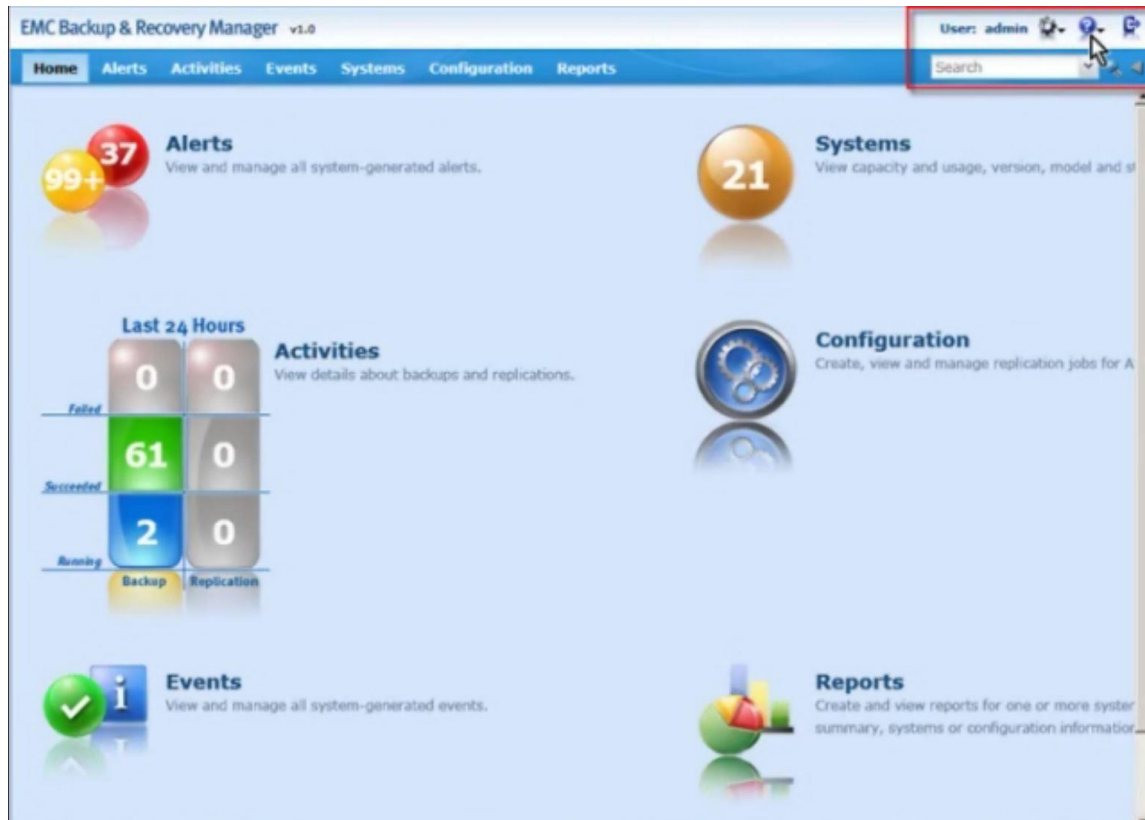
NetWorker does have its own reporting area in the NetWorker Management Console GUI view but it does not include as much performance information or information on other components in the overall backup solution. The DPA reports beyond the included reports in the NetWorker Management Console and DPA is especially helpful for those sites that are required to share reporting information for backup and recovery compliance.

## **Backup and Recovery Manager**

Backup and Recovery Manager software provides control and monitoring for NetWorker, Avamar, and Data Domain systems from a web browser or tablet. It requires no license and Android and iOS operating systems are supported. A list of supported web browsers and versions can be found in the NetWorker Software Compatibility Guide.<sup>4</sup>

Backup and Recovery Manager software runs on a virtual server and can collect alerts and information from NetWorker, Avamar, and Data Domain for monitoring purposes. In one view, it clearly displays the most important alerts and events that need to be addressed. The main page on the Backup and Recovery Manager user interface provides a one-panel style view of information from the three products with numbers to

indicate how many alerts and warnings are active, how many jobs have succeeded or failed in the last 24 hours, and an option to connect to a Reports sections to run reports to gather more detailed information.



**Example 1: of the Backup and Recovery User Interface** <sup>5</sup>

The main page display is broken into six areas: Alerts, Activities, Events, Systems, Configuration, and Reports.

1. The Alerts section allows a user to view alerts from all monitored systems including backup failures, errors, warnings, and media requests. Alerts often require user intervention.
2. Activities View shows detailed information on jobs that are running, completed, or queued on all monitored systems.
3. Events View shows all events generated by monitored systems in the enterprise. The ability to filter and acknowledge events is also provided.
4. The Systems View details information for all Avamar, NetWorker, and Data Domain systems in the enterprise.

5. The Configuration area allows the ability to configure basic Avamar replication.
6. Reports area allows a user to run backup summary, configuration report (lists policy objects), and system summary reports for all monitored systems in the enterprise.



Example 2: Backup and Recovery Manager Mobile Interface <sup>6</sup>



Example 3: Backup and Recovery Manager Mobile Interface <sup>7</sup>

The Backup and Recovery Manager architecture includes a Backup and Recovery Manager that uses REST API and a PostgreSQL database and adapters that are installed on the applications to be monitored. For NetWorker, the Backup and Recovery Manager Wizard can be used to install the adapter on the NetWorker Server. Starting with NetWorker 8.1, the install files for installing the adapter are included with the NetWorker server software but do need to be installed if Backup and Recovery Manager is going to be used.

## **EMC SourceOne**

EMC SourceOne (ES1) is an enterprise solution that can archive email, file, and Microsoft SharePoint servers. The content that is archived can be full-text indexed and searched to help meet regulatory, compliance, and legal requirements for archiving. EMC SourceOne can archive and organize content from Microsoft Exchange (Server and Office 365), IBM Lotus Domino, and SMTP mail servers.

### **Backing up SourceOne with NetWorker**

EMC SourceOne is a distributed system with components and data in more than one place. It is important that the system be put into a suspended state for backup and the different components are put into the suspended state in the correct order. Once the backup is complete, operations can resume, but again be started in the correct order.

Create a comprehensive list of the Microsoft SQL Server, Master, Worker, and Native Archive servers and the locations where information is stored by EMC SourceOne. Data locations need to include paths such as the Message Center, Full-text indexes, Native Archive, Logs, etc.

Discovery Manager server and its SQL database, the Supervisor Server and its SQL database, and the Microsoft SQL Reporting Services server may also need to be included in the backup groups in NetWorker.

The components that include Microsoft SQL databases such as the activity, archive, and search that are part of SourceOne will need to be backed up using the NetWorker Module for Microsoft (NMM). Discovery Manager, Supervisor (if installed), and Microsoft SQL Reporting Services Server (if installed) also have SQL databases that will need to

be backed up with NMM. With NMM, the MS-SQL data can be backed up without stopping the SQL Server services in a hot backup mode. The options used in NMM to back up these databases should follow standard database backup policies already being used within the organization. Please consult the MS-SQL database administrators for current strategies being used to perform online backups that will allow a point-in-time recovery of the databases.

The other storage locations need to be backed up as save sets to clients in NetWorker that are created for these hosts that own or can access the data such as the Message Center, Full-text indexes, Native Archive, Logs, etc. Depending on how SourceOne has been architected there may be multiple systems that will need to be installed with the NetWorker client software and backed up as clients in NetWorker client resource list and group resources. If the systems are virtual machines, NetWorker Protection for VMware or EBR backups may be an option which will simplify how many clients need to be created in NetWorker. Be sure to include them in the VMware Backup Policies. However, if the clients are virtual, an image backup and recovery would suffice for systems such as Master, Worker, and Native Archive servers. The MS-SQL data should still be backed up as an application through NMM and the Native Archive storage is likely external from the virtual environment so make sure it is accounted for.

Before backing up any of the systems or locations where SourceOne data is stored, it will need to be put into a suspended state for backup by using the EMC SourceOne backup scripts. The scripts can be found with the SourceOne installation software in the \Utility\Backup folder. Five vbs script can be found there to support the backup of SourceOne.

There are three suspend scripts and two resume scripts that can be run:

Filename of script	Description
ES1_ActivitySuspend.vbs	Suspends all workers prior to backup, does not allow searches.
ES1_ActivityBackupSuspend.vbs	Suspends all workers prior to backup, but allows searches.
ES1_NativeArchiveSuspend.vbs	Suspends the Native Archive archives prior to backup.
ES1_NativeArchiveResume.vbs	Resumes the Native Archive archives from a suspended state
ES1_ActivityResume.vbs	Resumes workers from a suspended state.

**Table 2: EMC SourceOne backup scripts<sup>8</sup>**

The Activity and Native Archive need to be suspended before the backup and then resumed after the backup. In NetWorker, a pre-script and post-script can run from the pre-command and post-command attributes in the client resource or by setting up savepnp backups. The savepnp command in NetWorker is used to call a command before the backup (pre) and a command after the backup (post).

To suspend activity for backup choose between ES1\_ActivitySuspend.vbs, which suspends activity including searches, or ES1\_ActivityBackupSuspend.vbs, which suspends activity but allows searches. For suspending the Native Archive, there is a ES1\_NativeArchiveSuspend.vbs script. Since there are two scripts to be run it will be necessary to put them together in one .bat file with 'nsr' or 'save' at the beginning of the file name and call that .bat file that has both suspend scripts in it.

For example, for a 64-bit version of Windows:

```
nsr_suspend_es1.bat
```

```
C:\Windows\SysWOW64\cscript ES1_ActivitySuspend.vbs
```

```
C:\Windows\SysWOW64\cscript ES1_NativeArchiveSuspend.vbs
```

For a resume .bat to be run after the backup

```
nsr_resume_es1.bat
```

```
C:\Windows\SysWOW64\cscript ES1_NativeArchiveResume.vbs
```

```
C:\Windows\SysWOW64\cscript ES1_ActivityResume.vbs
```

The ES1\_NativeArchiveResume.vbs script resumes the Native Archive archives from a suspended state.

If the SourceOne systems are distributed across enough different systems it may be necessary to schedule suspend and resume scripts in Windows Task Scheduler and have them run at a set time rather than rely on a pre-command/post-command on a single system. A NetWorker group can be configured to start after the time the suspend scripts are scheduled to run. If needed, the group could be started as a probe-based group with a probe script to run that will check that SourceOne is in the suspended state. The resume scripts can be scheduled through the Windows Task Scheduler as well to be run after all the SourceOne components are backed up.

Depending on where the Native Archive exists, during suspend/resume there may need to be a different backup or replication action that will need to be started to back up or copy the archive data. Exactly which strategy is used will depend largely on what type of storage the Native Archive files are on. Another consideration for shortening the amount of time the SourceOne environment is suspended is to consider using some kind of snapshot backup or snapshot with rollover so that the suspend - snap - resume can happen quickly and the copy to backup media (the rollover) can happen after the SourceOne environment is already started again. Options for managing the snapshot-type backup will depend on what underlying servers or storage the SourceOne systems have been installed. Some example features that could be used in NetWorker include



NetWorker Snapshot Management (NSM), NetWorker Snapshot Management for NAS (NSM for NAS), Block-Based Backup (BBB), NetWorker Protection for VMware (EBR), NetWorker Module for Microsoft (SQL VSS backup).

## **Conclusion**

This article does not propose to replace EMC published Installation Guides, Administration Guides, Release Notes, Integration Guides, etc. Many of these guides are written to help administrators implement and manage the strategies discussed in this article. While excellent resources, it can be difficult for a NetWorker Administrator to see a road map of where the pieces fit together by reading NetWorker's detailed product documentation. Many storage administrators are in charge of multiple products and roles making it nearly impossible to find time to review all of the guides before being thrust into the role of managing these environments. For example, I have met many customers that have assumed there were licenses or hardware to purchase additionally to their existing licenses to properly protect a technology of their clients when in fact, NetWorker was capable of doing it natively or by downloading a free add-on to NetWorker such as Backup and Recovery Manager. Knowing about options such as these may ease their ability to manage the environment.

This article can serve as a road map and overview of how NetWorker can coordinate protecting data in the enterprise with additional features and complimentary products that may be documented elsewhere, as opposed to being bundled in a single place or document.

## Appendix

1. EMC NetWorker Snapshot Management for NAS Devices Version 8.2 Integration Guide June, 2014 <http://support.emc.com>
2. EMC Avamar Version 7.1 Administration Guide January, 2015 <http://support.emc.com>
3. EMC Data Protection Advisor Software Compatibility Guide January, 2015 <http://support.emc.com>
4. EMC Backup & Recovery Manager Version 1.2 User Guide September, 2014
5. <http://support.emc.com>
6. EMC Backup and Recovery Manager 1.0 eLearning from EMC Education <http://education.emc.com>; this view very similar from version 1.0 to 1.2.
7. EMC Backup & Recovery Manager Version 1.2 User Guide September, 2014 <http://support.emc.com>
8. EMC Backup & Recovery Manager Version 1.2 User Guide September, 2014 <http://support.emc.com>
9. EMC SourceOne Email Management Version 7.1 Administration Guide September, 2013 <http://support.emc.com>

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.