



NETWORKER INTEGRATION FOR OPTIMAL PERFORMANCE



Mohamed Sohail
Shareef Bassiouny
EMC Corporation

Table of Contents

- Executive Summary 3
- NetWorker Integration with VMware..... 4
 - Agent-Based Backup..... 4
 - Proxy-based backups..... 4
- Avamar Integration.....14
 - NetWorker management of backups15
 - Replication of backups15
- The NetWorker with Avamar environment.....15
- Choosing deduplication clients16
- Deduplication best practices17
- Recovering Data23
 - Disaster recovery23
 - Causes of disaster24
 - Disaster recovery requirements24
 - Disaster recovery scenarios24
- Troubleshooting25
- NetWorker / Data Domain Integration26
- Achieving Data Domain’s best deduplication ratio – guidelines for top deduplication efficiency.27
- Appendix.....35
- Author Biography35

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

Executive Summary

This article discusses the sophistication of integrating multiple products—even products from the same vendor—to form a stable, consistent workflow, and the challenges that we meet to protect our data via an efficiently running backup solution. As EMC NetWorker® is one of the classic, sophisticated, and dominant backup solutions in current data center environments, its integration with different EMC or non-EMC products and how to run the setup smoothly and efficiently—with different devices and backup equipment without having a negative impact on the total performance—is an important topic of interest to many backup experts and IT professionals in general.

In this article, we illustrate many vital and important points in the integration of NetWorker and other products from EMC portfolio including Data Domain®, Avamar®, and VMware. The aim of this article is to enable a smooth integration between the different EMC BRS products and avoid being stuck on their limitations. It should also help the sales force provide customers a consistent view of EMC BRS products. Additionally, it will facilitate the mission of the Presales on differentiation points against other competitors, enabling them to sell a complete package and well-integrated solutions running in harmony as one product.

NetWorker Integration with VMware

This article summarizes what I have learned about VMware-NetWorker integration. I tried to make it FAQ-like with my main reference being NetWorker and VMware integration.

NetWorker supports two models for Virtual Machine backup and recovery. Additionally, it integrates with VMware vCenter server allowing a full view—through NetWorker Management Console—discovery of the Virtual Infrastructure (VI) ESX/ESXi hosts and the Virtual Machine guests running on those hosts. This feature—called Auto-discovery—enables the backup administrator to be aware of the latest changes made on the VI in terms of a Virtual Machines relation to their host ESX/ESXi host. This is an important part of running smooth backups for mission-critical Virtual Machines.

Virtual Machine (VM) backup has 2 models:

Agent-Based Backup: This is a classic model where a backup agent is installed on each VM, and the VM is treated as a physical host. This model is deprecated because of its impact on the Hypervisor resources (CPU, Memory, Network, etc.) within the time interval of the backup. Thus, backup policies should be defined to limit the number of simultaneous backup jobs that are running on each physical ESX Server, in order to limit the effect of backup jobs on the operation of other VMs. Despite its intensive resource requirement, it functions exactly as the classic Agent Based Backup for physical hosts in the sense that all procedures and limitations apply for backup and recovery.

Proxy-based backups: The modern model is proxy-based backup. In this model, a proxy host—where a backup agent is installed—communicates with the vCenter server controlling the Virtual Infrastructure and requests performing a Snapshot of the target VM. That Snapshot is then mounted through a transport mechanism and start the backup of that VM Snapshot. This mechanism offloads the backup effort from the Hypervisor—ESX/ESXi host—to the proxy host. The main advantage here is that the ESX resources are not depleted by the backup operation, thus preserving them for other production operations. As usual, new features come at the expense of additional configuration and operation procedures and with new limitations. Our aim is to clarify some of the most important points about NetWorker-vStorage API for Data Protection (VADP) integration.

Proxy-based Backup is not a new term in Backup and Recovery, it is known to be “running the backups of client A through another host B”; thus, the backup effort is offloaded from A to B. This is the case for VADP as the backup effort of the target VM is offloaded to the VADP proxy host.

Why do we want to do that?

Virtualization of production hosts aims to:

- preserve processing power of non-utilized physical hosts
- accelerate deployment procedures for new services
- optimize cost efficiency for buying new hardware

Backup operations are resource-intensive operations and thus they are normally executed during off-peak production hours; however, there are multiple situations where there is a lot to do even on non-business hours [re-indexing File systems, running analytics on production databases, running regular reporting scripts, and so on. In those situations, the data center operations team will have difficulty reserving time and processing bandwidth for backups. This is where PbB is a great solution by offloading backup effort to a proxy host.

How to Implement VADP backup using NetWorker

Simply prepare your proxy host; it should be able to communicate with the Virtual Center Server and the ESX servers. If you intend to use SAN transport mode, you will need to zone the Proxy host HBA adapters to your VMware DataStore LUNs. Install NetWorker Client on that host. The last thing you will need is to have a Virtual Center Server account, authorized for backup and recovery tasks execution. Refer to NetWorker and VMware integration guides for a list of rights required for the VADP user. Follow the configuration steps in the guide to be ready to run.

What are the Transport Modes? Are there pre-requisites to properly set up Transport Mode?

Transport Mode is the mechanism through which the proxy host mounts the VM snapshot in order to start executing the back up on the mounted snapshot. There are 4 transport modes:

- **Storage Area Network (SAN) mode:** In this mode, the Proxy host accesses the VM snapshot through the SAN. Of course, you will need to have the Proxy host HBA adapters zoned—just like the ESX servers—to the DataStore’s LUNs. Selecting this mode completely offloads the backup-related CPU, memory, and I/O load from the ESX/ESXi hosts to the

proxy host. The backup I/O is fully offloaded to the storage layer where the data is read directly from the SAN or iSCSI LUNs.

The VMs need to be hosted on either FibreChannel or iSCSI-based storage (ESX local DataStore will not function with this mode). The corresponding VMFS volumes—where the target VM is stored—must be visible to the VADP proxy host; Windows Disk Management snap-in should show those DataStore LUNs.

- **Hotadd mode:** Can the VADP proxy host be virtualized? Yes, it can! This is a good candidate for those who are not willing to dedicate a physical host for VADP proxy function and prefer to virtualize the VADP proxy or for those who need to rapidly deploy the VADP proxy without any changes on their SAN zoning. In this mode, the backup-related I/O happens internally through the ESX I/O stack using SCSI hot-add technology. This provides better backup I/O rates than NBD/NBDSSL. However, selecting this mode places backup-related CPU, memory, and I/O load on the ESX hosting the VADP proxy.

Hotadd mode requires a virtual proxy, and the ESX hosting the virtual proxy should have access to all the datastores where the VMs are hosted. So, if the datastores are SAN/iSCSI/NFS and if the ESX server where the VADP proxy resides is separate from the ESX server where the VMs are hosted, then:

- In the case of SAN LUNs, the ESX hosting the proxy and the ESX hosting the VMs should be part of the same fabric zones. This is a normal implementation step for VMware ESX servers, zoning them to the shared Datastore LUNs.
 - In the case of iSCSI LUNs, the ESX hosting the proxy and the ESX hosting the VMs should be configured for the same iSCSI-based storage targets.
 - In the case of NFS datastores, the ESX hosting the proxy and the ESX hosting the VMs should be configured for the same NFS mount points.
- **Network Block Device (NBD):** If SAN usage not an option for any reason—SAN access problem, design constraints, and so forth—then Network access may be the best option for the VADP operation. In this mode, the CPU, memory, and I/O load gets directly placed on the ESX hosting the production VMs, since the backup data has to move through that ESX host and reach the proxy over the network. NBD mode can be used either for physical or virtual proxy, and also supports all storage types. With the advent of Gigabit and 10GigE

network access technologies, this mode appeals to those who do not wish to use SAN transport.

- **Network Block Device with SSL (NBDSSL):** Encryption for network access security, NBDSSL transport mode is the same as NBD except that the data transferred over the network is encrypted. Data transfer in NBDSSL mode can therefore be slower and use more CPU due to the additional load on the VADP host from SLL encryption/decryption. So, this presents the usual performance versus security tradeoff.

Is there a best practice for Transport mode usage? How much does it scale in terms of concurrency? In other words, how many VMs can I back up concurrently/in parallel?

Yes, there are recommendations for usage of each transport mode. Here are few of them. A full list is available on the “Recommendations and considerations for transport mode” section of the NetWorker and VMware integration guide.

- **SAN transport mode:** This is the most scalable mode, supporting up to 50 concurrent backups when using backup to disk. This limit may grow to 100 when using DDBoost devices! However, DDBoost devices require a large amount of memory on the VADP proxy host (approx. 500MB per device).
- **Hotadd Transport mode:** A minimum of 4 vCPUs (≥ 2.66 GHz), with 8GB vRAM per virtual proxy. No more than 12 virtual disks should be mounted on the Virtual proxy at any given time. If any of your VMs has more than 12 Virtual disks, perhaps you should use NBD for that VM. Consequently, set the client parallelism of the VADP proxy client so that the virtual proxy will never mount more than 12 virtual disks concurrently. The ESX server hosting the virtual proxy must be running ESX 3.5 update 4 or later. If using DDBoost devices, size the Virtual proxy memory to 300MB per device.
- **NBD/NBDSSL Transport modes:** Make sure not to have more than 20 concurrent virtual disks mounted from any ESX/ESXi host. Set the client parallelism of the VADP proxy client to guarantee that limit.

Performance optimization rules include:

- Properly size the VADP proxy based on your parallelism requirements.
- Use multiple proxies if you have a large number of VMs or need a high parallelism.

- overall backup performance of VADP Proxy will be defined by the slowest component in the entire backup data path—Transport mode used, proxy resources (memory, CPU, Bus, etc.), and I/O load on the storage infrastructure at the time of the snapshot creation.

For more information refer to the “Performance optimization” section of the NetWorker and VMware integration guide.

Should I choose physical or virtual proxy? How will performance differ?

Whether to use a physical or virtual proxy should be determined based on performance requirements, the choice of backup targets, and available hardware. There are no observed performance differences between physical and virtual proxies when running on similar hardware.

Does Snapshot technology guarantee recovery of my data? What about file consistency of my database?

Performing a backup using VMware VADP creates a crash-consistent—the OS file systems consistency is guaranteed—snapshot of a virtual machine image. However, advanced VMware functionality allows a backup application using VADP to achieve application-level consistent backups. When performing a full VMware backup using VADP, in addition to VM quiescing—meaning freezing with data consistency guaranteed—vSphere version 4.1 and later provides application quiescing using VSS on Windows 2008 and later platforms. This functionality requires that the VMware tools package is installed on the VM guest. If VMware tools are not installed, there is no backup integration with the VSS framework and backups are considered only crash-consistent.

What if my application does not have a VSS writer, or if I am using Windows 2003, or any other problem?

Applications quiescing may be problematic in some cases. Issues on the virtual machine may prevent successful completion of quiescing VSS prior to snapshot creation:

<http://kb.vmware.com/selfservice/documentLink.do?externalID=1018194µsiteID=null>

This VMware article details this topic and provides a way to disable application level quiescing by removing the VSS driver component from VMware tools installation if it consistently causes snapshot failure. Another workaround for application level consistency is to stop/freeze the

application before the snapshot, and resume it after the snapshot through the Pre-freeze and Post-thaw scripts facility available in VMware tools. For further details about this facility, check <http://kb.vmware.com/kb/1006671>.

For more information, refer to the section “Advanced use and troubleshooting” of NetWorker and VMware integration guide.

Is there a best practice to guarantee snapshot success?

Some recommendations based on known problems:

- Schedule backups when very little I/O activity is expected on the virtual machine datastore, as this can impact the time required for taking the snapshot or removing the snapshot.
- It is recommended to keep at least 20% free space on all datastores for snapshot management. This is absolutely important to avoid multiple snapshot creation problems.
- In the case of VMs that have a large amount of change rate during backups, the snapshots can grow in size considerably while the backup is running. Therefore, ensure that the snapshot working directory on the VMFS datastore—which is, by default, the same directory where the VM configuration files reside—has enough space to accommodate the snapshot during the backup.
- Prior to VMFS 5, VMFS datastores could have different block sizes—chosen initially when the DS is created and its FS formatted. This block size parameter impacts the maximum allowed file size on the datastore. Starting with version 4.0, ESX and ESXi will compare the maximum size of a snapshot redolog file—which may grow to match the size of the virtual disk—with the maximum “allowed” size of files on the datastore where the snapshot will be created. If the file could grow beyond the maximum “allowed” size, ESX cancels the Create Snapshot operation and displays the error: *File is larger than the maximum size supported by datastore*. Thus, ensure that the snapshot working directory supports the size of all the disks attached to your target VM.

What about incremental backups? What are the supported backup types? Does NetWorker support FLR file level recovery? Does the backup type and level affect recovery?

There is a dependence on the OS of your VM. The default backup type for this Snapshot technology is the FULLVM, which is a full backup of the VM snapshot (also called Image level backup) and thus, no incremental or differential backup can be done here! However, NetWorker supports FLR backups only for Windows VMs (Linux FLR backups still under development); in fact, this is the default for Windows VMs. NetWorker will parse the disks of your Windows VM after mounting its snapshot and execute the backup level required.

This means that the VADP backup, where the saveset name is FullVM or *FULL* of Windows VM can be incremental or differential backups. This can be checked through mminfo output:

```
01/06/2013 04:45:32 AM 9099 MB 3773360628 cb full FULLVM
```

```
01/07/2013 04:41:00 AM 2741 KB 3689560678 cb incr FULLVM
```

Both saveset names are identical and both are FLR. They differ in the backup level.

It is important to note that: The incremental and non level-0 backups (differential) allow recovery of files (FLR through a NetWorker client installed on the target VM). However, Recovery of the full VM is only supported for level-0 *FULL* save set backups. In other words, the full VM recovery is a saveset recovery of the FULLVM saveset whose level is full.

What is Changed Block Tracking [CBT]? What is its advantage? Its prerequisites?

FLR backups consume a considerable amount of CPU and memory on the proxy host, because of the effort required to parse the File System (FS) of the VM. An enormous FS with millions of files may consume a large amount of the Proxy host resources and last for long periods of time. This parsing operation is required to execute all backup levels—as with any other file system backup—and cannot be avoided in full backups.

VMs running on ESX 4.0 or later hosts with Virtual Hardware version 7 can keep track of disk sectors that have changed. On a virtual machine, the virtual disk block changes are tracked from outside of the virtual machine in the virtualization layer; [ESX and vCenter]. When a backup is performed, NetWorker uses CBT to determine which files have changed since the last backup, and backs up only those files. This accelerates incremental backup processing by avoiding full file system parse. For more information, check <http://kb.vmware.com/kb/1020128>

Considerations and limitations about VADP environment setup and supported configurations

VMware tools must be installed on the VM to ensure that backups are taken in a consistent state. Also, VMWare tools are needed for client backup via FQDN/hostname.

For performance, EMC recommends using the VADP proxy host as the storage node. This provides the optimal configuration for any given transport mode as data transfer occurs directly from the ESX/ESXi datastore to the storage node.

The VADP proxy system must be running one of the following:

- Microsoft Windows 2003 (with at least SP1 installed)
- Microsoft Windows 2003 R2
- Microsoft Windows 2008
- Microsoft Windows 2008 R2

If you are planning to use tape drives as a backup target with a virtual proxy host, you cannot use Hotadd transport mode. VMware does not allow VMdirect with Hotadd. For more information, refer to “Support for tape drives in a VM” section of NetWorker and VMware integration guide.

Ensure that port 902 is open between your VADP proxy and the ESXi servers. For connectivity with vCenter, the TCP port is 443 by default.

For SAN and Hotadd Transport modes, it is best to disable automount on the proxy OS level to protect RDM LUNs from being automatically mounted to the VADP proxy when a VM with RDM disks is mounted onto the proxy. In some cases, this may lead to corrupting the RDM disk data. Refer to the “Diskpart utility for SAN and hotadd transport modes” section of NetWorker and VMware integration guide.

Optimal CPU load and performance when using DDBoost devices is observed with four concurrent backups per device. A lower number of parallel sessions to a single device does not achieve full performance while a higher number increases CPU load without additional performance gain. Based on the CPU load, there is typically no performance improvement from adding more than 3 DDBoost devices per proxy node.

When the datastore is almost out of space, VMware creates a snapshot named Consolidate Helper while attempting to delete snapshots. This snapshot cannot be automatically deleted by the backup application. To remove the Consolidated Helper snapshot, the VM must be shut down and the snapshot manually deleted from vCenter before the next backup. Otherwise, change files may accumulate on the datastore. The accumulation of such files can affect both the backup performance and the I/O performance of the virtual machine. Information about deleting the Consolidate Helper snapshot is provided in VMware knowledge base article: <http://kb.vmware.com/kb/1003302>. To avoid this annoying issue, ensure that there is always sufficient space available for snapshots.

VMs with physical compatibility RDM disks are not supported for VADP backups, because VM snapshots cannot be applied to such VMs. During NetWorker backup of a VM, no RDM-related information is backed up, and no RDM disks/data are restored upon VM recovery. You will need to make note of all LUNs that were zoned to your VM, and re-attach them to the VM after recovery. If you need to backup the data on the RDM disk, you will have to use the classic Agent based Backup (AbB).

You will not need to install NetWorker client on your Target VM for VADP backup. However, you will need to install NetWorker client for File Level recovery

Image level recovery is made only from a level full of the FULLVM saveset.

During the recovery of a full virtual machine (FULLVM save set), the recovered virtual machine will start in forceful powered off state because of a VADP snapshot limitation.

The following issue applies to NetWorker releases 7.6.2.1 build 638 and later. Traditional guest-based (client-based) backups are “not browsable” in the recovery GUI for VMs that are running a non NTFS filesystem and that have a mix of VADP and guest based backups. This issue does not apply to Windows VMs that are using NTFS. Additionally, save set recoveries are not affected and can be performed in the usual way.

When recovering a VM using SAN mode or Hotadd mode, make sure that the “SAN Policy=OnlineAll” and that the DataStore Target LUN/disk have its attribute “Read-only: No”.

VADP does not support IPv6, so you may need to disable IPv6 on your Windows proxy host.

It is recommended to keep the vCenter and VADP proxy as separate machines to avoid contention of CPU and memory resources.

An incremental backup of a VM is not supported after a hardware change, OS patch update, Service Pack update, drivers update, and so on. Perform a full image level backup after every change made at the operating system and hardware level on the VM.

Backup and recovery directly to a standalone ESX/ESXi host is not supported. The ESX/ESXi must be connected to either VirtualCenter or vCenter to perform backup and recovery operations.

There are limitations concerning VM disk Storage Configuration. The following VM disks are not supported and cannot be backed up using VADP:

- Virtual Machine OS containing GPT or dynamic disks [Windows]
- Virtual Machine OS containing partitions without drive letters [Windows]
- Virtual Machine OS containing uninitialized disks
- Virtual Machine OS containing unformatted partitions
- Virtual Machine configuration with Virtual IDE Disk Devices (only SCSI disks are supported)
- Virtual Machine configuration with independent disk mode. If such disks are detected during backup, they are skipped and a message is logged that indicates the disks were skipped. However, during an image level recovery, the disk is recovered without any data. If using independent persistent disks, you must use the traditional NetWorker-style backup for protecting the data on the independent persistent disks via the backup client installed inside the VM.

Limitations apply to non-English versions of the Windows operating system using vCenter for VADP. The following names should always contain only English characters:

- Backup VM display name in the left pane of vCenter
- Backup VM hostname/FQDN
- vCenter Datacenter name

- vCenter Resource pool name
- ESX datastore names containing the VM configuration files and virtual disks

VMs can only be restored to the same language OS vCenter that was used during backup. For example, you cannot recover a VM backed up from a Japanese OS vCenter onto an English OS vCenter.

VADP recovery can only be performed using the NetWorker User program. A command line recovery of the entire image will not work in cases where the backup was performed from a non-English vCenter.

On the machine where the VADP recovery is launched, the NetWorker package should be installed in English only without any language packages. You must unselect all the other language packages explicitly during the NetWorker installation.

Avamar Integration

EMC has integrated its widely deployed and well-established backup application, NetWorker, with Avamar[®] next-generation client deduplication. This integration effectively creates a unified platform to deliver next-generation backup and recovery alongside traditional approaches, e.g. tape, for those customers who want a more evolutionary approach to adopting data deduplication for disk-based backup.

Data deduplication has become mainstream technology for backup. In fact, the industry has nearly reached the point where the majority of large enterprises employ deduplication in their backup infrastructure. The capacity savings it delivers and the faster backups and restores it enables are simply too good to pass up.

EMC's integration of NetWorker and Avamar has produced a tightly coupled solution. By combining the NetWorker and Avamar clients into a single backup agent and managing Avamar within NetWorker resources, administrators can manage the entire backup environment through a familiar user interface and common workflow. This reduces complexity by simplifying the day-to-day management and further consolidating the backup infrastructure. Backup procedures do not have to change to adopt client deduplication. You also have one go-to vendor for backup implementation and support.

EMC is offering an evolutionary and pragmatic approach to inserting deduplication into your backup process. For enterprises that do not have the luxury of starting from a blank slate with their backup infrastructure, the ability to deploy traditional and deduplication backup using a single client agent and manage both from a single console offers an attractive solution. Consider it as your enterprise looks at how to adopt deduplication and back up to disk.

NetWorker management of backups

Deduplication save sets are handled by the integration in the same way as regular NetWorker save sets, except deduplication save sets consist of two parts that are stored separately. Both of these parts are tracked by the media index on the NetWorker server:

- The Avamar deduplication node stores the client backup data on the Avamar server.
- The backup metadata consists of a hash ID per file.

This information is used to identify which sub-file segments in a backup must be deduplicated or have already been deduplicated. The metadata is essential to recover the stored data back to its original non-deduplicated format.

Replication of backups

Backups that are stored on an Avamar deduplication node may be copied by Avamar procedures to a separate Avamar replication node for efficient disaster recovery or automatic failover of backups. NMC does not manage or track Avamar replication operations. Replication configurations and operations must be configured, launched, and monitored from the Avamar Console. Replication nodes must be created and configured with the assistance of EMC Customer Support.

The NetWorker with Avamar environment

The NetWorker with Avamar environment stores not only deduplicated client data but also the deduplication metadata and information that is necessary to recover the client data. Figure 1 shows the paths of the backup data and metadata that are processed as follows:

1. The NetWorker server initiates the backup of several client groups within its datazone.
2. The NetWorker storage node for the clients communicates with the clients to identify which data needs to be deduplicated and stored.
3. The backup data is deduplicated at each client and stored on an Avamar deduplication node located on the Avamar server.

4. The metadata that identifies the deduplicated information is stored on a supported device, such as an AFTD, on the NetWorker storage node.
5. The NetWorker server maintains tracking information for the backup operation in its media index, which is typically stored on a locally-attached disk or tape device.

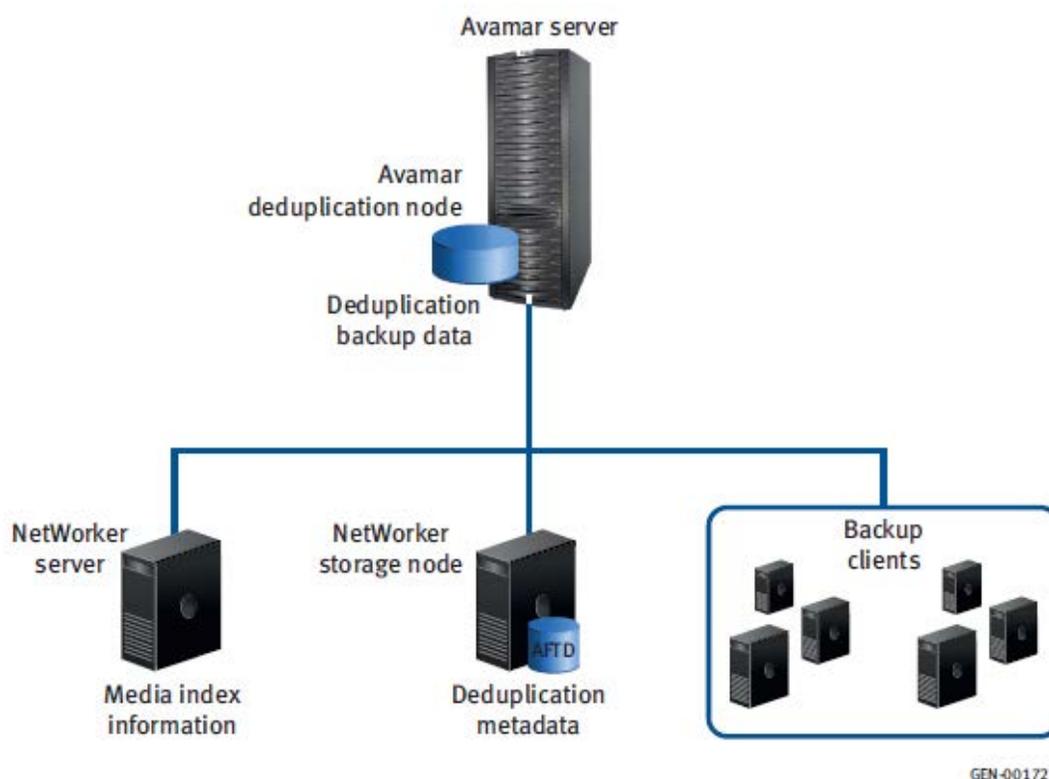


Figure 1

Choosing deduplication clients

You can decide which clients are good candidates for deduplication backups with the help of the following set of factors:

Include clients that can provide high deduplication ratios

To get the most efficient use of storage and backup schedules, choose clients that can provide high deduplication ratios. Deduplication ratio is the reduction in storage space that results from the deduplication/compression technology that is applied to save sets. Some backups may achieve ratios up to 500:1, but reductions of even 5:1 are extremely valuable.

Factors that impact deduplication ratio are:

- Change rate
- Retention period

- Data type
- Frequency of full backups

Exclude clients with low deduplication benefits

Certain performance factors can exclude clients from the benefits of deduplication backups.

These factors include:

- Number of client files
- Client application support
- Network bandwidth
- Client priority
- Client host performance

Faster backups versus faster recoveries

Your choice of backup schedule will depend on whether you give more importance to faster backups or faster recoveries. Table 1 compares estimated save set backup and recovery times relative to a non-deduplication full backup (x hours).

Backup type	Full initial backup	Subsequent daily full backups	Subsequent daily incremental backups (after a weekly full)	Data recovery time after 1 week
Non-deduplication	x hours	x hours	$0.1 x$ hours	$1.6 - 2.0 x$, if the full backup and all differential backups are to be restored.
Deduplication, daily full	$1.5 x$ hours	$0.2 x$ hours	n/a	Less than or equal to x (only one backup need be restored).
Deduplication, daily differential	$1.5 x$ hours	n/a	$0.05 x$ hours	$1.5 - 2.0 x$, if the full backup and all differential backups are to be restored.

Table 1

Deduplication best practices

The performance of Avamar deduplication nodes depends on good scheduling practices, attention to the size and configuration of the Avamar server and—most importantly—Name resolution of Avamar server, nodes, and clients. The network environment has an impact on

hostname resolution methods and you should follow the recommendations of the operating system vendor:

- The deduplication client must be able to resolve the names of the NetWorker server and the configured Avamar deduplication node.
- Avoid renaming the Avamar server, deduplication nodes, or deduplication clients.

These operations should be done with assistance from EMC Technical Support:

Note: Renaming the Avamar server, or renaming deduplication or replication nodes that are hosting unexpired backups will adversely affect recovery of the existing save sets.

Note: Renaming clients that have unexpired deduplication backups will adversely affect data recovery.

Backup and replication schedules

Schedule backups and replications so that they do not extend into the Avamar maintenance blackout window or the NetWorker maintenance window:

- At the start of the blackout window, running backups have a 15 minute grace period before they are cancelled. If this cancellation produces a partial backup, it will be stored on the Avamar server for seven days to facilitate the next backup of the client.
- The NetWorker software will retry interrupted backups up to the value set for client backup retries.
- Schedule and perform replication operations only at times when backup operations will not be impaired, typically at the end of the backup window.
- A recovery from a replication node may be started only after all of the replicated data is stored and on the node.

Extending your retention policies

If you extend a client's retention policy, carefully monitor the Avamar server's capacity usage. The Avamar server can fill up much faster than expected.

Rollbacks to checkpoints

The NetWorker software is unaware of changes that result from rollbacks to checkpoints on the Avamar system. Rollbacks to checkpoints can result in lost backups, lost clients, and lost deletions.

The following situations can occur:

- Avamar deduplication save sets that were created between the checkpoint and the start of the rollback become unrecoverable.
- Deduplication clients created in the Avamar domain during the rollback period become lost.
- Save sets that were deleted during the rollback period become undeleted. Provide EMC Customer Support with the `/nsr/logs/nsravamar.raw` log file if you need assistance to list the deletions on the Avamar server.

Avamar server workloads

Other operations performed with the Avamar server can impact performance. The Avamar server may be shared among different NetWorker datazones as well as the native Avamar backup software. However, added complexity impacts server performance, scaling, and storage space usage. Use EMC Data Protection Advisor (DPA) to manage complex environments.

- Use of the Avamar server for other operations that are not available through the NetWorker Management Console can impair performance and should be done only with the assistance of EMC Customer Support.

Metadata storage requirements

The backup metadata consists of a hash ID per file backed up to the Avamar server. This information is used to identify which sub-file segments in a backup must be deduplicated or have already been deduplicated. The metadata is essential to recover the stored data back to its original non-deduplicated format.

The metadata also enables all backup, clone copy, restore, disaster recovery, monitoring, and reporting operations to be managed and tracked through the NMC display.

Metadata device

Clone the metadata device volume or save sets to protect this essential information and ensure disaster recovery.

- The metadata device should be dedicated for deduplication clients and not used by other types of clients, pools, or data.
- The volume for the metadata device should be labeled and belong to a deduplication pool.
- Metadata device names should be simple.

- The device path and name should be unique for each NetWorker datazone.
- The name of the device volume should not exceed 60 characters.
- Locate each device within its own directory even if there is a dedicated file system or partition.
- Use a clear and consistent naming convention if other similar devices are used.
- The metadata device should have enough space to accommodate all the full and differential backups for the retention periods involved.

Depending on the file characteristics, each metadata file could generate between 160 bytes to 1 KB of metadata. Thus, 200 million files could require as much as 200 GB of metadata on the metadata device.

Reclaiming expired storage space

When save sets on a deduplication node exceed their NetWorker retention policies, the NetWorker server begins the process of deleting the expired data and recycling the storage memory. Either on startup or every six hours, NetWorker invokes a temporary maintenance process to start the deletion of expired data on the Avamar deduplication node and the associated metadata on the NetWorker storage node.

- If you delete the NetWorker server's media index entry or overwrite the volume label, this action will only free up space that the deduplicated save set used.
- If an Avamar server is removed from the NetWorker datazone, the internal Avamar retention policy reverts to "no expiration" and saveset data will be retained on the Avamar server.

The deduplication node cannot be deleted in NetWorker as long as there are expired save sets in the NetWorker queue to be deleted on the Avamar server.

The deletion of expired data from the Avamar server should be monitored with NMC or the Avamar Console to ensure that this process runs successfully and that the server does not fill up with redundant data, especially if Avamar replication is being used.

Replicating backup data

To ensure added data protection of backup data—for disaster recovery or failover scenarios, for example—save sets and backup metadata stored at the primary location can be replicated to a secondary location. This process involves two separate procedures:

1. Avamar replication of backup data: The backup data, stored on a deduplication node on the primary Avamar server, is replicated to a replication node on an Avamar server at a secondary location. Single save sets or the entire volume of the deduplication node may be replicated.
2. NetWorker clone of backup metadata: The backup metadata, stored on the primary NetWorker storage node is cloned to a NetWorker storage node at a secondary location. By default, the cloned metadata retains the original NetWorker browse and retention policies, but these can be changed to give the replicated data different policies. Staging (moving) of the metadata is not recommended.

Figure 2 illustrates a basic disaster recovery environment. In this scenario, a data loss at the primary site can be recovered by restoring data from the secondary site. A failover scenario, in which the entire primary site fails and the secondary site continues backup and recovery operations for the clients would require a NetWorker server (not shown) at the secondary site with media index information cloned from the primary NetWorker server.

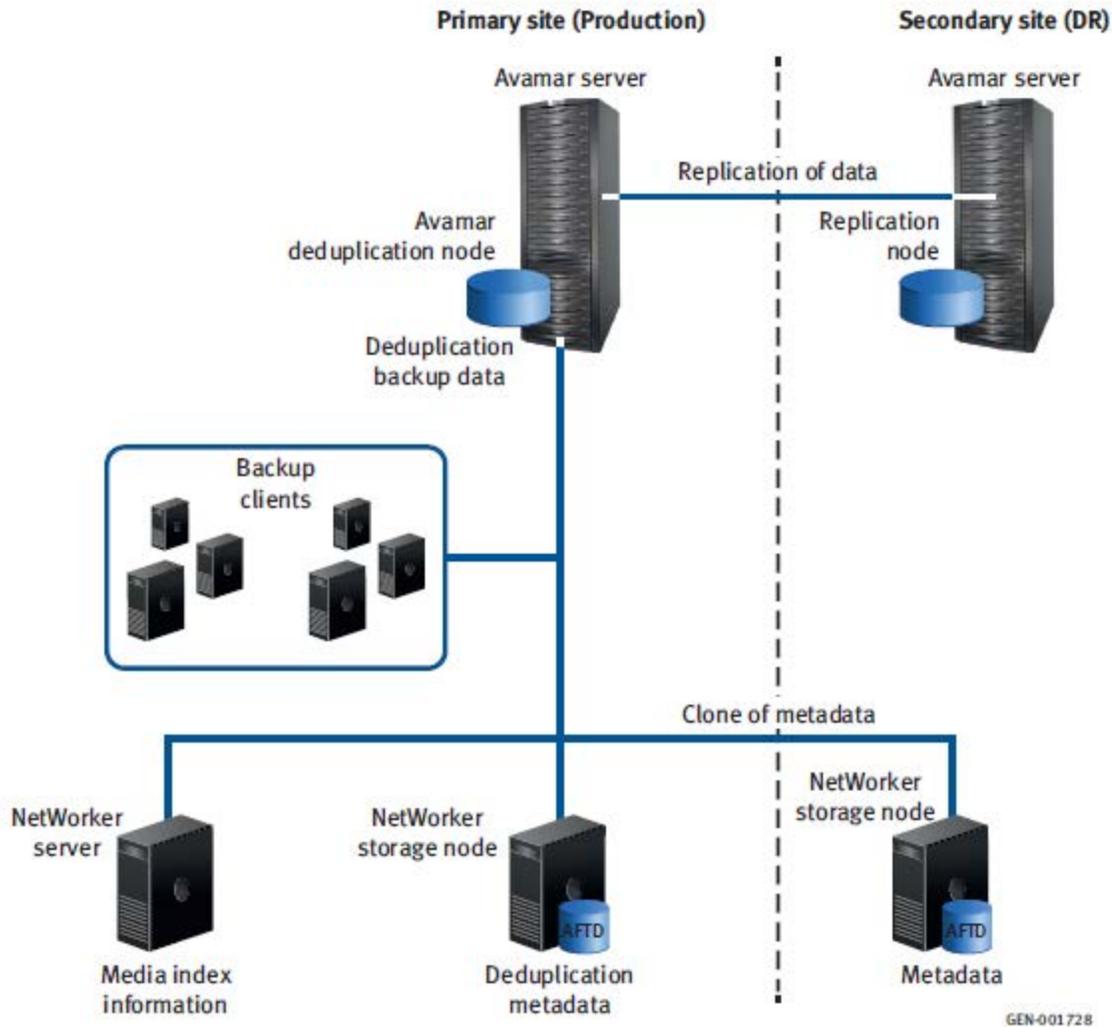


Figure 2

Replication and clone requirements

The following requirements must be met before replication of Avamar deduplicated backup data and clone of the corresponding metadata can be performed:

- Replication nodes should be created and already exist before integration with the NetWorker server, preferably when the Avamar server is installed.
- Data replication must be configured, scheduled, and monitored on the Avamar server. The NetWorker software cannot perform these operations.

- An Avamar replication node must also be an Avamar deduplication node and must be part of the NetWorker server's datazone, even if it is at a geographically distant location.
- The NetWorker storage node, which will store a clone copy of the metadata for the replicated backups, must also be part of the NetWorker server's datazone.
- A replication node does not require a dedicated host. The host may also serve a deduplication node for different clients and backup groups.
- Schedule or manually start replication operations at times that will ensure backup operations are not impaired, typically at the end of the backup window. A replication started immediately after a backup finishes will experience a short delay.
- All of the replicated data must be available on the replication node before a recovery operation can be started from that node.

Recovering Data

Provided that the NetWorker and Avamar software have been properly configured with a deduplication node or replication node, the data recovery from either of the nodes is basically the same as for a NetWorker storage node, except that the data is stored in two components in different places:

- Deduplicated client backup data is stored on an Avamar deduplication node, which is on an Avamar server.
- Backup metadata, which is the hash information used to deduplicate and track the stored backup data, is stored on a NetWorker storage node. This metadata is required to revert the deduplicated data to its original non-deduplicated format.

Recovery requirements

The requirements for recovery from a deduplication node or replication node are as follows:

- All of the deduplicated or replicated data must be available on the deduplication or replication node before the data can be recovered.
- Both the Avamar deduplication node and the NetWorker storage node must be online during the recovery of deduplicated data.

Disaster recovery

A disaster is any loss of data in which the computing environment required to restore that data is not available. Disaster recovery (DR) is necessary when ordinary data recovery procedures

are not sufficient to recover the computing environment and its data to normal day-to-day operations.

Causes of disaster

A disaster can result from any of the following situations:

- Hardware or debilitating software failures
- Computer viruses that corrupt the computing system
- Infrastructure interruptions, inconsistencies, or loss of services, such as communications or network connections that result in damage to the computing environment.

Disaster recovery requirements

A complete DR environment provides a secondary site with systems that copy all of the information involved in each completed backup at the primary site.

The two sites may be configured to provide disaster recovery for each other, with each serving as both a primary and secondary site with different datazones for different clients.

DR requires maintenance of the following systems:

- DR Avamar server with deduplicated client data replicated from the primary Avamar Server
- DR NetWorker storage node with deduplication metadata cloned from the primary NetWorker storage node
- DR NetWorker server with media indexes cloned from the primary NetWorker server.

Disaster recovery scenarios

Primary Avamar server is lost, secondary Avamar server becomes primary

In this scenario, the primary NetWorker server and storage node survive the disaster, but the primary Avamar server is lost. A DR Avamar server at the secondary site contains backup data that was replicated over a WAN or IP connection from the primary site.

To recover from this scenario:

- Reconfigure the NetWorker server to promote the secondary Avamar server to become the primary deduplication node.

- Manually configure the NetWorker Client resources for the new deduplication node. The new primary Avamar server at the DR site will store client backups over the WAN or IP connection.
- After the primary site is fully restored and enabled for backup operations, demote and resynchronize the secondary site.

Entire primary site is lost and service is restored at the secondary (DR) site.

In this scenario, the primary site has entirely lost its NetWorker server, NetWorker storage node, and the primary Avamar node. The secondary Avamar server is available at the disaster recovery (DR) site, which will become the new primary site.

To make the DR site the new primary site:

1. Install the NetWorker software on the system at the new site.
2. Configure the NetWorker server with the storage devices and the surviving (secondary) Avamar node.
3. Enable the devices with the No Mount option.
4. Run the scanner -B device_path command and note the bootstrap SSID information.
5. Run the mmrecov utility with the bootstrap information to recover the NetWorker server:
6. After the recovery is complete, shut down the NetWorker services.
7. Rename the res folder to res.old and res.R to res.
8. Restart the NetWorker services.

Troubleshooting

The following sections will help identify and resolve common configuration and operation issues.

Restore operation fails

A restore fails. For example, a disaster recovery with Windows OSSR may fail to recover “critical volumes.”

To resolve, ensure that the client has access to the Avamar server.

Avamar server fills with expired save sets that were not deleted

Expired save sets might not be removed from the Avamar server. In some cases the server fills to capacity and requires assistance from EMC Customer Support to resolve. The following situations can contribute to this problem.

Replication node is not properly configured in NetWorker

The Avamar replication node configured in NetWorker must use the FQDN of the Avamar/destination_directory:/REPLICATE/node_name.

Requested deletions become out of sync

Save sets on the NetWorker deduplication node can become out of sync with the save sets on the replication node.

For example, NetWorker sends a request to delete a save set and while the deletion is completed on the replication node, a blackout window prevents the deletion on the NetWorker deduplication node. When the next replication is performed, the save set is recopied from the NetWorker deduplication node to the replication node. This save set will never be deleted from the replication node.

Until this synchronization problem is resolved, a workaround can be used. Install a crontab script that replays the deletions on the replication node three days back in time.

NetWorker / Data Domain Integration

Most of the information here is gathered from personal experience but in order to assure that I am not writing my own subjective view I tried to align—as much as possible—with the most important document on the topic; “NetWorker Data Domain integration guide”.

NetWorker is a backup system, thus it needs to have a data storage target. Classically, the backup storage media was tape volumes, with the evolution of the data center requirements the most important feature required by the data owners. Consequently, performance—how high a data rate can the storage technology reach for reading and writing—is the most important feature to storage administrators. LTO technology— from the good old LTO2 up to the current LTO5—became quite common in today’s data centers with added features such as hardware-enabled compression and encryption.

However, for business-critical applications and tight backup windows, more performance is required. Along with advances in the storage device industry and the cost reductions achieved, Backup to Disk (B2D) became the next wave of backup data storage. B2D is simply enabling the backup software to use disk storage for backup data either through the Virtual Tape Library (VTL)—which is a virtualization engine that makes your backup system think that it is connected to a real tape library—or through the backup system vendor development of a backup data

storage format on plain hard disks, resulting in much faster backup data storage and thus narrowing the backup window and accelerate data restore. For long-term storage, the backup data is staged from the high performance disk storage to the less performing, yet much more economical, tape storage, which is still the choice for long-term retention and archiving. This model has been known as “The Data Life Cycle”.

Explosive data growth makes disk space requirements for B2D quite large, negatively impacting its cost effectiveness. Data deduplication came on the scene, allowing storage of more data using significantly fewer hard-drives.

Data Domain is a deduplication storage device that performs quite well, either target-based, where data deduplication takes place on the target storage or source-based deduplication, where data deduplication takes place partially on the source host. Its integration with NetWorker is maturing with development of new features and functions to enhance performance.

Achieving Data Domain’s best deduplication ratio – guidelines for top deduplication efficiency

Deduplication Ratio represents the reduction in storage space that results from the data deduplication/compression technology. Ratios of 20:1 are considered to be broadly achievable and reductions of even 5:1 are extremely valuable.

Important factors that contribute to the deduplication ratio include:

- **retention periods:** Ratio increases with longer data retention periods. The longer the stored save sets are retained, the greater the chances that identical data already exists in storage that can be used to deduplicate each subsequent backup
- **type of data being backed up:** Some types of user data such as text documents, slide presentations, and spreadsheets are known to contain redundant data and are good deduplication candidates. Other types of data such as audio, video, and scanned images already consist of compressed data which may not yield high ratios due to the fact that compression is a type of redundancy elimination],
- **data change rate:** Data that does not change much between backups (low change rate) produces high deduplication ratios as data chunks already stored will not be stored again. Typically, the first full deduplication backup of these data types yield low reductions, but subsequent backups will typically produce high deduplication ratio.

How much will be gained by integrating NetWorker with Data Domain?

The answer depends on how you want to run your backups and your priorities; of course, the most accurate calculations will require a talk with your pre-sales consultant, but an interesting tool may be Asset Management and Planning (AMP), a virtual appliance designed to enable EMC customers to self-assess their usage of certain EMC software assets. It is a free, self-installable, virtual appliance that can be downloaded from Powerlink and installed on the customer's ESX server. Using AMP is as simple as defining discovery policies and setting parameters for the desired frequency of data collection—everything else is automated. Customers will have the ability to view, print, and export resulting usage reports.

You may also consider the capacity-based licensing model which enables all features of NetWorker and licenses the capacity of the protected data, thus removing a great deal of licensing complications inherent in NetWorker's classic licensing model.

I already have DD working with NetWorker through classic configuration (B2D file devices through CIFS or NFS, on other DataZones, where I am running a relatively old NetWorker version). I use the VTL. What will the DDBoost feature do for me? What are its pre-requisites?

DDBoost enables managing fewer devices that support high concurrency (parallel ingestions of data from multiple sources). The current device structure is quite dependent on many storage devices (NAS or SAN for B2D) and numerous VTL devices for environments that prefer VTL. The number of devices that you will need with DDBoost is much less making the structure and management of your backup system setup much simpler and easier, without any changes in performance. In fact, for operations such as cloning savesets between DD devices, there is a considerable performance gain.

Pre-requisites are simply a currently supported version of NetWorker (7.6 or later, if you plan to upgrade from an older version). However, for the best performance, I recommend NetWorker version 8 which will enable you to use the latest DDBoost Library. On the Data Domain side, you will need at least DDOS version 4.8, but I strongly recommend having DDOS 5.2; again, for best performance, unless you have a certain problem that prevents you from running the latest versions. For more information about upgrading, please refer to Appendix : Upgrading from a NetWorker 7.6 SP1 release on "NetWorker Data Domain integration guide"

How does DDBoost accomplish that? How does it provide source side deduplication?

DDBoost is a library that is integrated into the NetWorker client and Storage node software package. Through this integration, it participates in data deduplication on the NetWorker Storage node side (sometimes on the client side as well, through the version 8 new feature, client direct with DDBoost) to optimize network bandwidth usage, which optimizes overall performance.

The DD Boost software consists of two components:

1. The distributed segment processing (DSP) component, which reviews the data that is already stored on the Data Domain system and sends only unique data for storage.
2. The DD Boost library API, which enables the NetWorker software to communicate with the Data Domain system.

Does this you mean that I may be missing something with my NetWorker 7.6?

Yes I believe so. While running the latest version of Data Domain DD Boost embedded library solves many problems and provides top performance, the most appealing feature added in Version 8 is the backup workflow: namely Client Direct.

Client Direct, also known as direct file access (DFA), is a NetWorker feature that enables clients with IP network access to the Data Domain [or a backup to disk device] system to send backup data directly to the B2D device or DD Boost storage devices, bypassing the NetWorker storage node. The storage node manages the devices used by the backup clients but does not handle the backup data. The Client Direct feature is used by default when it is available (when the client has access to the DD device).

The Client Direct feature leverages the DD Boost distributed segment processing (DSP) software component that is installed as part of the NetWorker client software, version 8.0 and later. During backup, the DSP software on the client deduplicates the backup data before the Client Direct component sends the deduplicated data directly to the DD Boost devices. By working together, the DD Boost and Client Direct features can provide highly efficient data deduplication, transmission, and storage for multiple concurrent client backup operations. Backup bottlenecks are removed from the storage node, and network bandwidth is better utilized. Figure 3 depicts the difference between the two data flow paths.

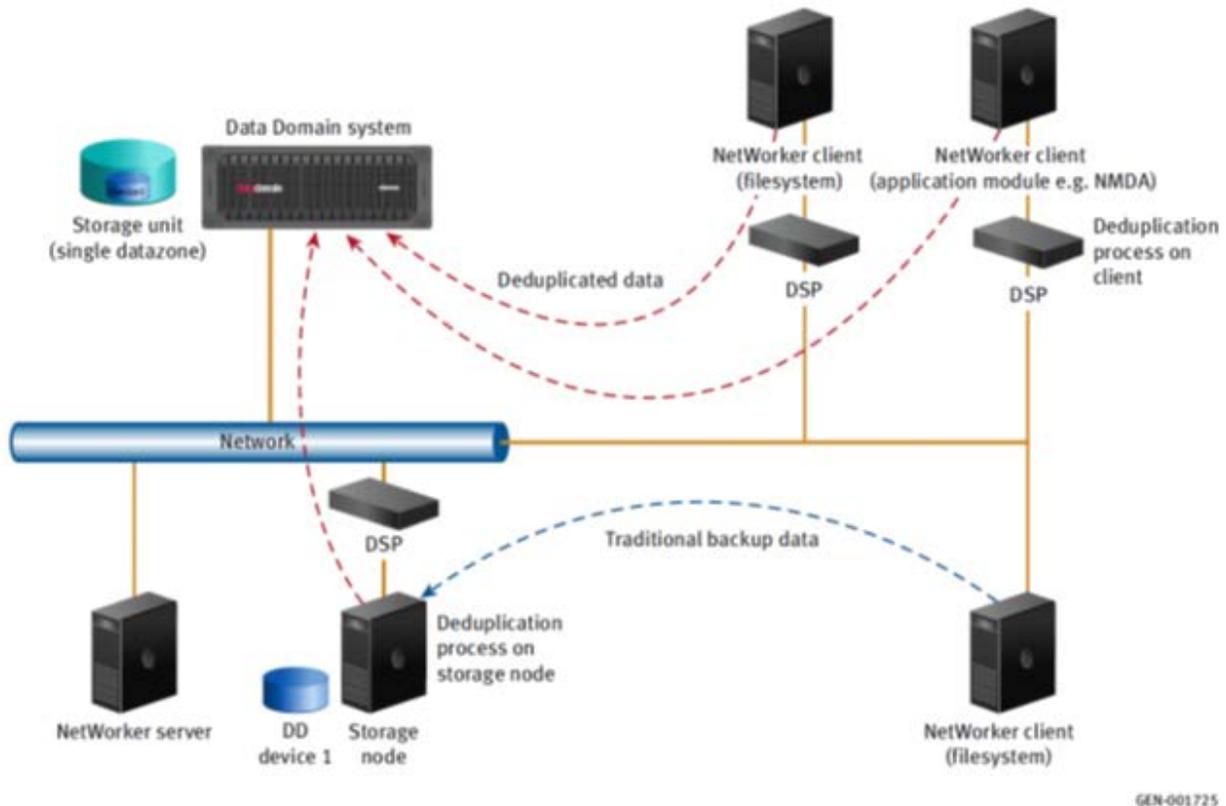


Figure 3

For more details, please check the “Client Direct deduplication environment” section in the NetWorker Data Domain integration guide.

What are the pre-requisites for using DDBoost Devices? How many resources do they require on the host side?

The answer depends on the usage—how many parallel streams from client side; if the client is using DFA with DDBoost; if no DFA, how many clients will be sending their Data simultaneously to DDBoost enabled Storage node; and so on. However, there are some indexes.

- On the client side, DD Boost clients require a minimum of 4 GB of RAM at the time of backup to ensure optimum performance for Client Direct backups.
- On the Storage node side, allowing for other types of devices and services on a typical storage node, a storage node should have a minimum of 8 GB of RAM if hosting DD Boost devices. Each DD Boost device requires an initial 24 MB of RAM on the storage node and Client Direct client. Each DD Boost save session requires an additional 24 MB. The default max sessions of 60 sessions per DD Boost device requires 24 + 1440 MB.

How does using a DDBoost device enhance cloning?

For disaster recovery purposes, and sometimes for archiving purposes, cloning backup data is required. Cloning backup data stored on Data Domain requires converting the data format from their deduplicated format stored on Data Domain to their native non-deduplicated format. This process of Data Reading from Data Domain consumes time and effort, but is necessary if we are cloning to any traditional backup media (tape or disk).

However, if we are cloning from Data Domain to another Data Domain device (DR site cloning, for example), we can optimize the operation by using the clone-controlled replication (CCR) or “optimized clone” when the source and target devices are DD Boost devices. This happens through maintaining the deduplicated format of the data through replication between the source and target devices, so when NetWorker detects that source and target devices are DD Boost devices, it will run the cloning operation as a replication for the required savesets between the source and target DD Boost device.

The advantage is that the cloning process in this case will use minimal resources—processing effort, network bandwidth, time, and storage—as it replicates only the new chunks from source to target devices. Of course, the initial data transfer from source to target is like a first full backup (initial seeding of deduplicated data chunks, which requires extra effort than normal B2D because of the required chunking, hashing, and compression efforts), with the exception that the data is transferred/replicated in deduplicated format. On subsequent clones, only new chunks are replicated. This greatly optimizes normal clone workflow, which typically requires a full recovery of the data, then writing all the data to the target device/media.

How does the DD Boost device look on the Data Domain side in terms of the storage structure? Is there any recommendation on the number of devices per DD host?

DD Boost devices storage is structured as follows:

- Storage on the Data Domain system is based on storage unit (SU) parent folders that contain the DD Boost devices. These parent folders are also called managed trees (MTrees). By default, each SU handles a single NetWorker datazone and is named after the short hostname of the NetWorker server.
- Data Domain operating systems 5.0 and 5.1 should use no more than 14 active SUs.
- Each DD Boost storage device is identified by a SU child sub-folder and is associated with a single NetWorker storage volume.

- There is no recommended limit on the number of DD Boost devices that you can create, but for best performance, use fewer devices and more backup sessions on each device.

What are Data Domain Archiver systems? Is there anything special to do on the NetWorker side when integrating with such systems?

A Data Domain Archiver system is a specific model designed for the management of long-term data retention. NetWorker DD Boost devices may be created on a Data Domain Archiver system that runs DD OS 5.0 or later software.

On Data Domain Archiver systems, data can be stored in two tiers of storage units; the active tier and the archive tier

All incoming data is first placed in the active tier file system, similar to standard Data Domain systems. This tier is used for short-term data storage and recovery. The active tier can be used for backup provided appropriate data movement and retention policies are used for the devices—NetWorker-side configuration of Retention policies. The creation of separate SUs for backup data is recommended.

Data on the active tier can subsequently be moved to the archive tier based on data movement policies applied at the SU level—Data Domain-side configuration of Storage Unit retention and data movement policies from active tier to archive tier.

Each storage unit has a single data movement policy that applies to all of the devices it serves within its corresponding NetWorker datazone, whether they are NetWorker AFTD or DD Boost devices.

How would the Data Movement policies within a Data Domain Archiver affect NetWorker?

A Data Domain SU data movement policy determines when data is moved from devices in the active tier to devices in the archive tier. Policies should typically not move data to the archive tier for short periods, less than 30 days retention. The policy and movements are internal to the Data Domain Archiver system; NetWorker software has no awareness of them.

In order to make NetWorker leverage this feature, you should make it execute the data movement between storage tiers; which means, create devices in SUs on the archive tier and use NetWorker Staging / cloning to move the data between those tiers. You can also leverage CCR performance by using DD Boost devices in your Storage tiers, then use CCR to clone data

between the different devices in each tier. It is recommended to replicate between two different SUs for CCR on the same Data Domain system that includes the Extended Retention Software feature, so that different retention policies can be applied and the data can be managed efficiently.

Finally, some consideration and limitations

- NetWorker 8.0 provides both write and read functionality on B2D devices and, consequently, new DD Boost devices. This is unlike earlier releases that used separate read-only DD Boost mirror devices for restore operations. Installing NetWorker 8.0 and later removes legacy read-only DD Boost mirror devices. This advance has a backward fallback penalty; after upgrading to NetWorker 8.0 or later from prior NetWorker releases, a downgrade to the earlier NetWorker release is not supported. After the downgrade, DD Boost devices created with NetWorker 8.0 or later will be unavailable and legacy devices and data will require manual reconstruction (scanning all devices that received data after the upgrade).
- Due to the DD Boost library compatibility, Data Domain systems that run DD OS 5.0 or later software support DD Boost devices that run NetWorker 7.6 SP2 or later software.
- Data Movement policies between Active Tier devices and Archive Tier devices is internal to Data Domain Archiver; NetWorker has absolutely no awareness about it.
- The following is an important checklist for proper connectivity between various entities:
 - The storage nodes and all Client Direct clients must have network access to the Data Domain system where the data will be stored and recovered from.
 - If the NetWorker server and the NMC server are located on different systems, then both require access to the Data Domain system for administration and monitoring purposes.
 - NetWorker support for DD Boost devices does not distinguish network types (LAN, WAN, or MAN) and can successfully operate where packet loss is strictly 0% and latency is less than 20 milliseconds.
 - Currently, the DD Boost devices support only an IP network and do not support SAN (fiber channel) data transport.

- Ensure that DNS and hosts file are consistent, and that DNS response time is within acceptable ranges.
 - You may like to leverage Data Domain interface grouping [ifgroup] to secure Network connectivity against NIC failures and consolidate NICs bandwidth for performance optimization.
 - If a firewall exists on any network path, make sure that required ports are open bi-directionally. The required ports are listed in the Firewall requirements section of the NetWorker Data Domain integration guide.
- For DFA backups, make sure that the “Checkpoint restart” is not enabled on your client configuration. If this feature is enabled, the backup will revert to the classic Storage node-based data flow path.

Appendix

- VMware integration: EMC NetWorker/EMC Data Domain Deduplication Devices integration guide”, P/N 300-999-722
- DataDomain integration: NetWorker/VMware integration guide” P/N 300-999-729
- Avamar integration: NetWorker/VMware integration guide” P/N 300-013-563

All are included in the Documentation portfolio

https://support.emc.com/docu41460_NetWorker_8.0_Documentation_Portfolio.pdf?language=en_US

Author Biography

Mohamed Sohail

Global Account System Engineer

Mohamed is a Global Account System Engineer in TSSO Organization at EMC.

Mohamed has been working in the IT industry for more than 8 years, 3 of them with EMC. Formerly with the Oracle support center in Oracle Egypt and a technical trainer at Microsoft, Mohamed holds a BSC in Computer Science from Sapienza University, Rome, Italy and a B.A in Italian from Ain Shams University, Egypt.

Mohamed is a certified EMC Proven Professional - Backup, Recovery Solutions (BRS) and VNX Specialist.

Shareef Bassiouny

Technical Support Engineer – BRS / NetWorker

A Backup Recovery NetWorker Specialist, Shareef is a Technical Support engineer in the GTS Organization at EMC.

Shareef has over 12 years IT experience in operations, implementation, and support; over 3 of those spent with EMC NetWorker support. Shareef holds a Bsc. in Telecommunication Engineering from Cairo University. His previous role was leading a Dedicated IT Customer Support Desk that handled Data Center Operation and Change Management at Orange Business Services.

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.