

Backup and Recovery for VMware Using EMC Next-Generation Backup Solutions

Best Practices Planning Guide

Abstract

VMware® virtualization provides the power to consolidate servers and automate processes, enabling IT-as-a-service, but it can come at the cost of extra storage, backup resources and administrative challenges. EMC® next-generation backup solutions address these challenges by reducing redundant data across VMware data backups while operating at high speed; providing cost-effective replication for fast disaster recovery (DR) and centralized monitoring and management of all data protection operations. This white paper reviews best practices planning for a backup, recovery and a DR approach for VMware using EMC Data Domain deduplication storage systems with EMC Avamar and EMC NetWorker backup software that integrate with VMware vStorage APIs for Data Protection (VADP).

January 2012

Copyright © 2012 EMC Corporation. All Rights Reserved.

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

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

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

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

VMware, vCenter, VMware Consolidated Backup, vMotion, VMware High Availability, VMware Distributed Resource Scheduler and vSphere vStorage API for Data Protection (VADP) are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number h9569

Table of Contents

Executive Summary	5
Challenges of Backup and Recovery for VMware Environments.....	5
Introduction	6
Audience.....	6
EMC Backup and Recovery Solutions – Overview.....	7
Integrated Architecture.....	7
EMC Backup and Recovery Solutions for VMware.....	10
Image-level Backup	11
EMC Avamar – Image Level Backup	11
EMC NetWorker – Image Level Backup.....	12
EMC Data Domain Integration.....	13
Application Consistent Backup	14
EMC Avamar – Application Consistent Backup	14
EMC Avamar Integration with EMC NetWorker.....	15
EMC Avamar integration with Data Domain.....	16
EMC Data Domain Boost.....	16
EMC NetWorker – Application Consistent Backup.....	17
EMC NetWorker Module for Microsoft Applications.....	17
EMC NetWorker Module for Databases and Applications	18
Long-Term Retention of Backups.....	19
EMC Data Domain Archiver.....	19
Centralizing Management of Backups	20
Conclusion.....	22
Appendix	23
Backup Application Software Option Settings Recommended for use with Data Domain systems.....	23
Additional Resources	24

List of Figures

Figure 1: EMC Integrated Backup and Recovery Solutions	7
Figure 2: Avamar Client-Side Deduplication Moves Less Data	8
Figure 3: EMC NetWorker	9
Figure 4: EMC Data Domain Deduplication Storage Systems	10
Figure 5: Avamar Image-Level Backup	11
Figure 6: Avamar Proxy Server Load Balancing	12
Figure 7: NetWorker - VMware VADP Integration	13
Figure 8: VADP Enabled Backup with DR Workflow – Conceptual Diagram.....	14
Figure 9: Avamar agents inside of Guest OS	15
Figure 10: NetWorker with Avamar.....	15
Figure 11: Avamar Integration with Data Domain.....	16
Figure 12: EMC Data Domain Boost	17
Figure 13: NetWorker - Guest-Level Backup	18
Figure 14: NetWorker Module for Databases and Applications	18
Figure 15: Data Domain Archiver	19
Figure 16: Avamar Centralized Management	20
Figure 17: NetWorker - Unified Data Protection.....	21
Figure 18: DPA – Gain Complete Visibility of your Data Protection Environment	22
Figure 19: EMC Data Protection Advisor.....	22

Executive Summary

VMware® is the global leader in server virtualization for cloud infrastructure. Liberating virtual systems from physical hardware constraints enables server consolidation and increases host utilization, resulting in reduced hardware and environmental costs (heating, cooling and datacenter floor real-estate) supporting green initiatives. Simplified infrastructure management and automation reduce operational expenses for competitive advantage.

Challenges of Backup and Recovery for VMware Environments

With all of the benefits of server virtualization come some new challenges. Automating the installation and deployment of servers can lead to the proliferation of virtual servers, which could introduce bottlenecks with traditional backup and recovery processes that write to physical tape drives or plain-disk storage targets. The journey towards virtualization can stall when backup and recovery limits the number of virtual machines that can co-exist on a given ESX host.

When protecting data in VMware environments, restore requirements traditionally have dictated performing both full virtual machine (VMDK) image backups, and guest OS level backups for operational file recovery and application consistency. When backing up both VMDK images and application files inside of the guest OS, the resulting backup data is typically highly redundant. Without deduplication, legacy tape or conventional disk backup storage targets require a lot more storage to protect a virtual environment. Traditional operational backup and recovery methods limit the journey towards virtualization for the following reasons:

- **Application servers are slow during backup and the backups take too long**
 - Backing up virtual machines directly with traditional backup agents impacts performance on the application servers and can bottleneck the VMware / ESX host's shared resources (CPU, memory, network interfaces)
 - Application consistency for structured and semi-structured data (i.e. databases and email messaging systems) must be maintained without bottlenecking the ESX host's shared resources
- **Too much backup media is being consumed**
 - Separate backups for DR, operational file recovery and application consistency not only over run the backup window, but they consume more target backup media
- **Compliance is being compromised**
 - Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs) service levels are adversely affected, jeopardizing compliance

EMC® backup and recovery solutions address these challenges through the power of advanced deduplication. These solutions help accelerate virtualization adoption by increasing the number of virtual machines that can be effectively managed on an ESX host. EMC leads the market¹ in the integration of deduplication software and systems, providing unique and powerful high-speed, deduplication as a foundation for next-generation backup and recovery architectures. This industry leading deduplication technology extends throughout the backup solution to provide simple, network-efficient replication for offsite disaster recovery as a more reliable and cost-effective alternative to tape. Powerful centralized management and reporting, an integral part of the solution, provides the operational control and transparency, ensuring predictable data recovery.

Introduction

The purpose of this whitepaper is to exemplify how EMC backup and recovery solutions provide powerful and unique next-generation data protection for challenging VMware environments. With the solutions illustrated in this paper, you will have the knowledge needed to engage your local EMC account team; to assess your specific environment and determine which of EMC's integrated backup and recovery products can be utilized to architect solutions that best serve your unique data protection requirements.

Audience

This white paper is intended as a guide for backup administrators, systems engineers and partners to understand EMC's efficient next generation backup solutions for the VMware virtualization platform. It describes how EMC backup and recovery deduplication solutions simplify and reduce the cost of backup and recovery of virtual machines and resident data. A basic understanding of virtualization platforms and experience with backup and recovery software and concepts is assumed.

¹ IDC Study – Worldwide Purpose Built Backup Appliances http://www.emc.com/collateral/analyst-reports/11530-idc-ww-pbba-2011-2015-forecast.pdf?cmp=nlc-dedupenews-august-idcshowswhy&AM_ID=4998589&ActID=27626&M=89c5ccf3-143f-4272-ae99-d61e4305e999

EMC Backup and Recovery Solutions – Overview

The key to protecting data inside of VMware virtual environments is data deduplication. Additionally, centralized management and operational control of the backup and recovery infrastructure plays an integral part in successfully protecting the environment.

Integrated Architecture

EMC backup and recovery solutions are integrated to enable flexibility and a variety of options in solving particular backup challenges. Pain points can be addressed independently and ultimately solutions can blend into a cohesive backup and recovery strategy. For example, Avamar software can manage writing, replicating and recovering data to and from the Avamar Data Store, or a Data Domain system through easy configuration wizards and built-in monitoring and reporting facilities. Likewise, NetWorker can manage Avamar and Data Domain backup targets, in addition to legacy tape devices and disk array snapshots.

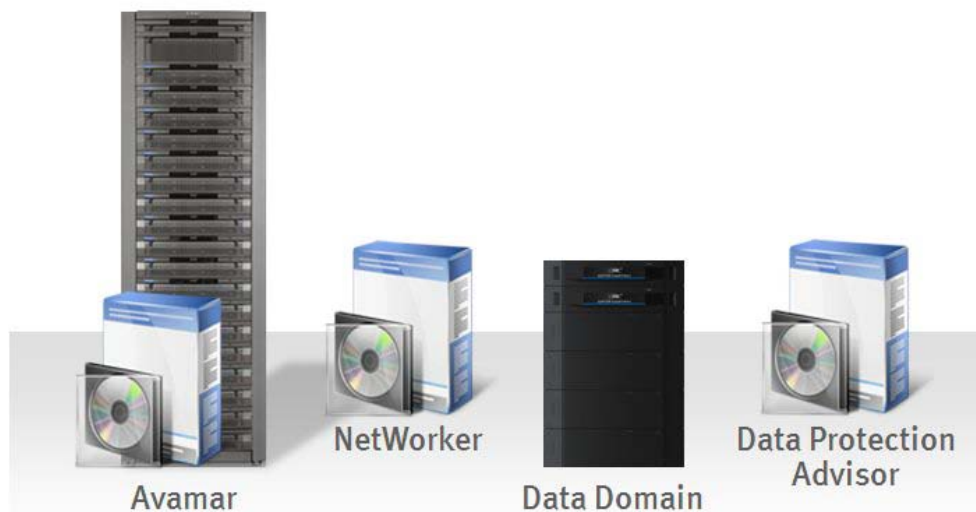


Figure 1: EMC Integrated Backup and Recovery Solutions

EMC Avamar Deduplication Backup Software and Systems

EMC Avamar is an integrated backup and recovery solution with software and hardware providing faster backups, reduced network traffic and simple single-step restores. Avamar’s unique client-side global deduplication technology eliminates redundant backup data before it is sent over the network and stored. As an enterprise class backup and recovery solution, Avamar is optimized for both physical and virtual environments – enabling a smooth transition during every stage of the virtualization journey. The Avamar Data Store offers high availability and reliability with its, “Redundant Array of Independent Nodes” (RAIN) architecture, enabling daily recoverability verification and DR via replication. Avamar is also tightly integrated with Data Domain systems for increased scalability and solution flexibility. Without

introducing management overhead, its management console controls all scheduling, policies, retention and recoveries for both the Avamar Data Store and Data Domain systems.

Avamar is optimized for VMware backup and recovery and is the industry's leading choice. It provides a variety of options, including image-level backup and guest-level backup. A simple one-step recovery is as easy as point and click to retrieve a file, folder, or database. Tight integration with VMware vCenter Server provides ease of use for managing the backup and recovery environment from a "single pane of glass". One key benefit is the ability to quickly see from within the backup and recovery management framework which VMs are protected and, more importantly, which ones are not. Avamar has been tested and certified with VCE Vblock Infrastructure Platforms in the EMC test labs providing assurance and proven reliability for administrators of these large scale virtualized environments.

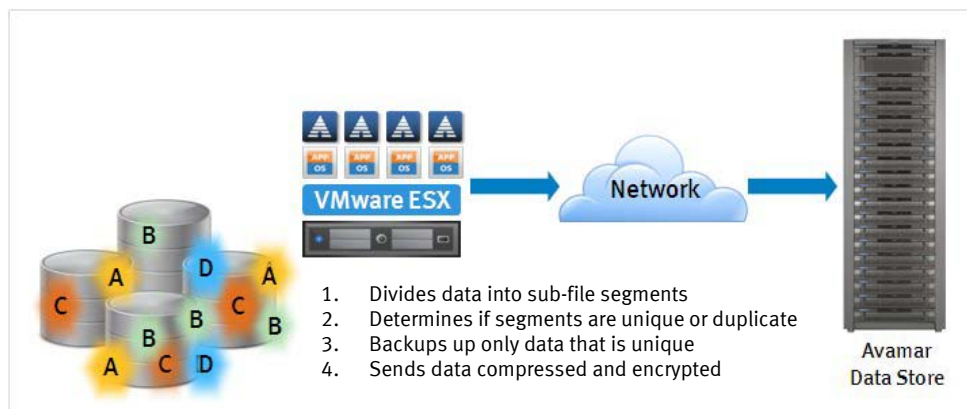


Figure 2: Avamar Client-Side Deduplication Moves Less Data

EMC NetWorker Unified Backup and Recovery Software

EMC NetWorker is a powerful application that provides unified control of the entire backup and recovery environment. It delivers a single management interface, necessary to provide greater simplification and easier management to meet the full range of existing and new requirements that are now challenging most organizations.

NetWorker is designed to support many different client and device types, and to provide one operational process and discipline to move and manage the data from primary storage to an onsite backup copy as well as cloning the data offsite for DR.

NetWorker is also integrated with Data Domain systems and Avamar. Avamar deduplication is built right into the NetWorker backup client software. These solutions drive new efficiencies in backup and recovery for VMware.

- Single solution with industry-leading deduplication
 - Avamar
 - Data Domain
- Integrated software and hardware
 - Simplicity
 - Predictability
 - Supportability

File systems and applications




Figure 3: EMC NetWorker

The combination of NetWorker and Data Domain allows for centralized, automated, and accelerated backup and recovery across the enterprise. Data Domain systems are ideal for use with NetWorker backup-to-disk capabilities as an alternative to tape. It can be deployed quickly and fits into current backup workflows and policies.

In addition, Data Domain Boost (DD Boost) software provides advanced integration between NetWorker and Data Domain. DD Boost distributes parts of the deduplication process to the NetWorker storage node or application client to provide the following benefits:

- Up to 50% faster backups
- Reduced LAN bandwidth required
- Reduced CPU utilization on NetWorker server
- Clone-controlled replication
- Automated configuration
- Monitoring and reporting

EMC Data Domain Deduplication Storage Systems

EMC Data Domain deduplication storage systems reduce the amount of disk storage needed to retain and protect data by ratios of 10-30x and greater, making disk a cost-effective alternative to tape. These systems can protect multiple petabytes of logical capacity in a single system enabling customers to retain data online and onsite for longer retention periods, as well as providing faster and more reliable restores. With throughput up to 26.3 TB/hour, Data Domain systems allow more backups to complete sooner while putting less pressure on limited backup windows.

EMC Data Domain Replicator software transfers only the deduplicated and compressed unique changes across any IP network, requiring a fraction of the bandwidth, time, and cost, compared to traditional replication methods, for faster “Time-to-DR” readiness.

Data Domain systems are qualified with all leading enterprise backup software and archiving applications and easily integrate into existing environments.

Data Domain systems come in a large range of sizes and performance levels to support any virtual server environment. Usable capacity starts at just over 1.6 TB and scales to hundreds of terabytes.



- **Scalable Systems**
Reduce storage by an average of 10-30x and protect petabytes of logical capacity in a single system
- **High-speed inline deduplication**
Complete backups faster
- **Network-efficient replication**
Send only deduplicated and compressed data across the network
- **Easy Integration**
Qualified with all leading enterprise backup and archiving applications

Figure 4: EMC Data Domain Deduplication Storage Systems

All Data Domain systems benefit from the EMC Data Domain Data Invulnerability Architecture – continuous verification, continuous fault detection and healing, and other resiliency features transparent to the backup application. For more information on Data Domain Invulnerability Architecture, review this white paper at <http://www.emc.com/collateral/software/white-papers/h7219-data-domain-data-invul-arch-wp.pdf>

EMC Backup and Recovery Solutions for VMware

Image-level backup for disaster recovery and operational file recovery is available for VMware vSphere environments with the integration of the vStorage API for Data Protection (VADP) into the backup application software, which provides for centralized backup of virtual machines, reducing ESX host resources on production application systems by executing off-host, LAN free backup processing and provides for flexible backup windows.

EMC backup and recovery solutions are engineered to manage VMware image-level backups with either Avamar or NetWorker backup software and can be integrated with a Data Domain system.

Intelligent Avamar agents or NetWorker client agents provide the tools needed to manage application consistency without adversely bottlenecking ESX host resources.

Image-level Backup

EMC Avamar – Image Level Backup

EMC Avamar takes advantage of VMware's vStorage APIs for Data Protection (VADP) to perform efficient image-level backups. The VADP approach involves creating a VMware snapshot of a virtual machine, presenting it to a proxy host – which is run in another virtual machine – then backing up the VMDK file utilizing that separate proxy host. This approach is preferred because it moves the backup process away from the virtual machine and likewise ESX resources where the applications are running. This is the optimal way to avoid any performance impact to the application VM itself. In this case, you would deploy the Avamar proxy host, treat it as the backup client, and send backups directly to the Avamar Data Store or Data Domain storage system.

- Changed Block Tracking support
 - Scan & copy the least amount of data to backup and recover entire virtual machines
- Unique automated virtual proxy management
 - Load balance image backups across multiple proxy servers
- Flexible image recovery options
 - Restore individual files from image backups (Windows)
 - Recover to original, existing or new virtual machines



Figure 5: Avamar Image-Level Backup

Because of VMware advancements and APIs, Avamar includes the ability to run the proxy server as a virtual machine allowing for VMware to dynamically “hot-add” volumes to that proxy without moving data across the physical or virtual network. Utilizing this mechanism Avamar is able to back up large virtual machines in just minutes (something to think about when considering backup for large virtual computing environments with hundreds of virtual machines to protect such as VCE's Vblock).

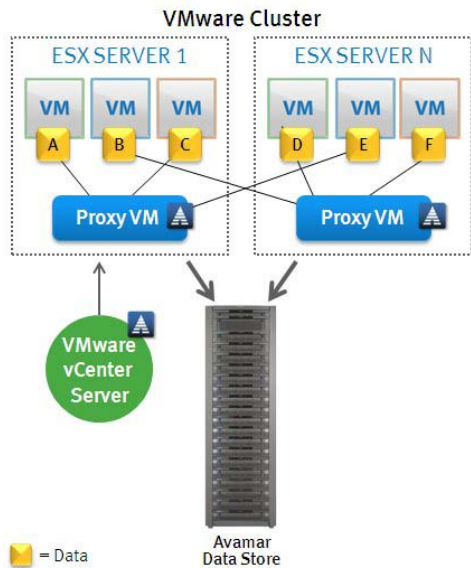


Figure 6: Avamar Proxy Server Load Balancing

Avamar can automatically load balance jobs across virtual proxy servers to assure the most efficient performance during the backup window. Jobs can also be directed to available proxies to handle high-priority workloads. This automated approach helps drive down costs by having fewer proxies to manage.

Avamar can also leverage advanced capabilities provided by VMware’s VADP APIs to only back up changed blocks. This lessens the scanning process that Avamar has to do, and speeds up the backups even further. Avamar’s architecture was designed to reconcile daily changed block backups into full instances for recovery. This is different from most legacy backup applications that require at least a periodic full backup. This means Avamar can process the changed blocks identified by VMware into a new recoverable, full image on an ongoing basis. As a result, recovery is much faster and more reliable. With Avamar 6.0 and greater, an even faster restore process has been developed. Leveraging the CBT mechanism on the restore, Avamar tracks the delta of changed blocks to speed the recovery process by restoring only the unique data segments required to restore the VM backup image requested – resulting in accelerated restores.

One feature that companies really would like is the ability to do one backup, and have the option of a bare-metal restore or a more granular file-level recovery. Taking advantage of its integrated file system, Avamar can provide this capability, without mounting the VMDK in a separate location to do so.

Finally, just as in physical environments, Avamar’s deduplicated replication of protected VMDK files allows for network-efficient disaster recovery – and importantly, great flexibility in how you can restore – either to the original virtual machine or to a new location.

EMC NetWorker – Image Level Backup

NetWorker 7.6 SP2 and above; delivers integration with the VMware vStorage API for Data Protection. VADP provides the facility to centrally and non-disruptively backup VMware vSphere hosts from a central proxy server.

During backups, NetWorker asks VMware through the API to create a snapshot of virtual machine data directly from the VMware datastore. The snapshot is then mounted and the data is copied directly to backup target media.

Through VADP, a single backup is able to deliver both a virtual machine image-level as well as individual file-level recovery. This saves time by eliminating the need to back up two different ways to achieve the desired level of recovery.

- Centralized backup for vSphere hosts
- Image and file-level recovery with a single backup
- Physical and virtual proxy support
- Deduplication support
 - Avamar deduplication on proxy
 - Storage node with DD Boost on proxy




Figure 7: NetWorker - VMware VADP Integration

With VADP, the backup proxy can be either a physical or virtual server. This allows the flexibility to choose based on individual requirements. NetWorker also integrates with the Changed Block Tracking (CBT) feature of the VMware API for incremental backups.

Either Avamar or Data Domain can be used in conjunction with NetWorker in this configuration. The NetWorker client featuring Avamar deduplication can be installed and used on the proxy, or the data can be sent to a Data Domain system via NFS or Data Domain Boost.

EMC Data Domain Integration

Leveraging a Data Domain system and a backup application integrated with VADP, removes the complexity and custom scripting previously required and makes configuring backup and DR workflows largely transparent to the backup administrator. Figure 8 depicts an overview of the actual steps involved in the workflow.

Note: The Data Domain system depicted in Figure 8 is configured with network-efficient replication to a remote site for DR.

1. Backup agent with VADP connects with vCenter, flushes source disk volumes with VSS², then creates a VMware vStorage VMFS³ snapshot of virtual machines on shared storage
2. Snapshots are mounted on backup application server
3. Unique data segments are copied to local Data Domain system target
4. Unique data segments are replicated to remote Data Domain system for DR

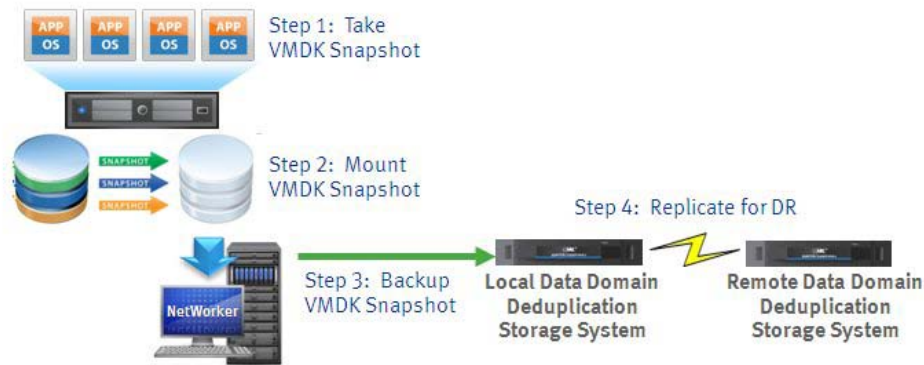


Figure 8: VADP Enabled Backup with DR Workflow – Conceptual Diagram

Application Consistent Backup

Many companies today continue to choose to backup VMware at the guest level. The main advantage to this approach is the level of flexibility and features provided for application consistency requirements. This approach is also great for file-level backup and recovery across multiple server and OS platforms.

EMC Avamar – Application Consistent Backup

For guest-OS level backup, the Avamar agent is installed directly on the virtual machine host operating system. When the backup runs, data is reduced within the virtual machine, and only new unique segments are copied and stored. The key benefit is that the backup bottleneck caused by contention for shared resources is alleviated. Avamar decidedly helps companies optimize consolidation ratios and get the most out of their VMware investment.

By running Avamar application-specific agents from within the virtual machine, application consistency can be guaranteed. Since Avamar employs a capacity based licensing scheme, you are only charged for the deduplicated capacity that you consume rather than the number of agents you deploy. This licensing method is especially great for virtual environments where virtual machines tend to come and go.

² For supported Microsoft[®] Windows[®] operating systems

³ Access the following link for more information on VMware vStorage VMFS
<http://www.vmware.com/products/vmfs/overview.html>




Figure 9: Avamar agents inside of Guest OS

- Application-specific agent resides within each VM ensuring application consistent state for hot backups
- Deduplicates within and across virtual machines
- Reduces resource contention and speeds up backups
- Provides file-level restore for Windows, Linux and Solaris
- Control cost with zero-dollar agents

EMC Avamar Integration with EMC NetWorker

With NetWorker, the common management of NetWorker and Avamar clients allows deduplication to be defined through standard NetWorker backup client configurations in the NetWorker Management Console. This centralized management of client agents is available to increase the effectiveness of backup for file systems and application modules in both physical and virtual environments. This includes support for Microsoft Exchange, SQL Server, SharePoint, and Hyper-V as well as Oracle, SAP, DB2, Informix, Lotus Notes, and Sybase.

The ability to manage NetWorker and Avamar from a centralized management console is ideal, especially for existing implementations, because it does not force the administrator to learn a different workflow or GUI.

- Supports file systems and applications
- Physical and virtual
- Manage traditional and deduplication backup
 - Single client or application module
 - Common administration
 - Full indexing and browse
 - Common recovery

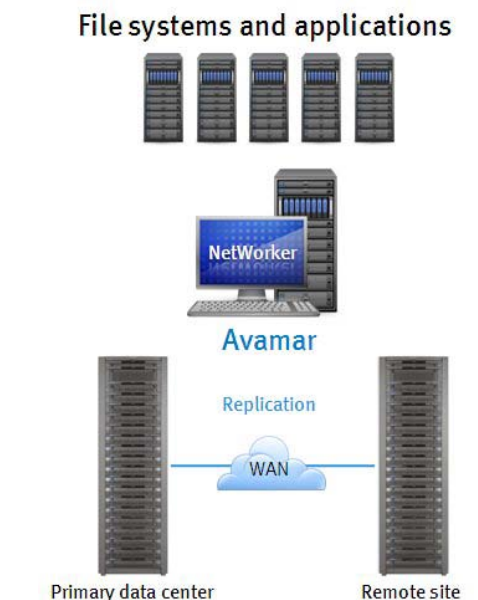


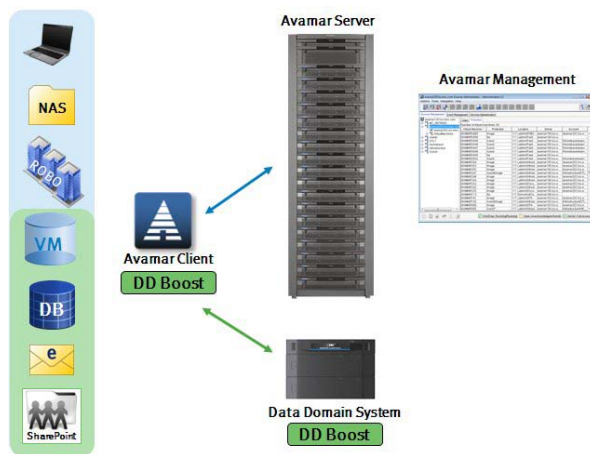
Figure 10: NetWorker with Avamar

The NetWorker Management console simplifies the management of the deduplication operations. Results of client backups with or without deduplication are available through common NetWorker Management Console reporting and monitoring tools. Recovery of any client and any data is possible through common recovery interfaces – graphical or command line.

EMC Avamar integration with Data Domain

The integration of Avamar with Data Domain provides the ability to selectively perform high-speed, scalable backups for specific data types, from an Avamar client to a Data Domain system. This combines two industry leading deduplication technologies into one seamless, integrated solution.

Allowing Avamar to move workloads to Data Domain takes advantage of Data Domain’s industry leading performance and scale, while leveraging Avamar’s simplicity and efficiency of management. Specifically, Oracle, Microsoft Exchange, SharePoint, SQL Server, and VMware image backup can be sent to Data Domain. Other types of data can continue to be directed to Avamar storage. This gives you the flexibility to select the best deduplicated backup approach based on your various application workloads. Typically, large files with high daily change rates are best targeted to Data Domain, while most other typical files are best targeted to Avamar. Regardless of the storage target, simplified, centralized management provides the operational control needed for smooth operations and compliance.



- High-speed scalable Avamar backups to Data Domain systems
- Broadens Avamar enterprise usage
- Solves more customer problems
- Direct selected backups to optimal storage systems based on workload attributes
- Provides access to an additional 285TB of usable capacity

Figure 11: Avamar Integration with Data Domain

EMC Data Domain Boost

EMC Data Domain Boost (DD Boost) distributes parts of the deduplication process from the Data Domain system to the backup server or application client. In addition to storage node support, NetWorker 7.6 SP2 or later supports DD Boost-based backup from application hosts for Microsoft applications and databases. This is driving new efficiency for users with NetWorker and Data Domain. By sending only unique data from the NetWorker server or application client to the Data Domain system, less LAN bandwidth is required, backups are 50 percent faster and the whole aggregate system more manageable.

DD Boost⁴ dramatically increases the aggregate throughput, up to 26.3 TB/hour, and reduces the amount of data transferred over the network by 80 to 99 percent. These efficiencies can help eliminate future costs by leveraging existing backup servers and Ethernet networks.

- **Increases backup speed by up to 100%**
- **Reduces network traffic**
- **Clone-controlled replication**
 - Schedules replication
 - Catalog awareness
- **Ease of use**
 - Wizard automated configuration
 - Monitoring and reporting




Figure 12: EMC Data Domain Boost

With DD Boost, backup applications can control replication between multiple Data Domain systems and provide backup administrators with a single point of management for tracking all backups and duplicate copies. This paradigm allows backup administrators to efficiently create DR copies of their backups over the WAN using DD Replicator software and keep catalog consistency for easy disaster recovery. This also provides the flexibility for administrators to manage different retention periods for each copy of data.

With NetWorker, the Data Domain replication process is managed by standard NetWorker cloning, ensuring that NetWorker can recognize and manage a replicated (remote) copy of data and assign unique retention policies to it. The administrator has the ability to schedule the cloning process to run at a time that is most appropriate for the business.

EMC NetWorker – Application Consistent Backup

EMC NetWorker enables network-efficient application consistent backups by integrating EMC Data Domain Boost with the EMC NetWorker Module for Microsoft Applications (NMM) and the EMC NetWorker Module for Databases and Applications (NMDA).

EMC NetWorker Module for Microsoft Applications

EMC NetWorker Module for Microsoft Applications can be used inside the guest OS to achieve application-consistent and restartable backups for virtual applications such as Microsoft SQL Server, Exchange, SharePoint, Active Directory and Hyper-V. With guest-level backup, the NetWorker client can be installed on each VMware guest and protect each virtual machine – just like protection is achieved for physical machines. This is generally very easy to do, and for companies who want commonality in installation and workflows, this can be a valuable option.

⁴ For more information on DD Boost, please reference the DD Boost data sheet at <http://www.emc.com/collateral/software/data-sheet/h7034-datadomain-boost-sw-ds.pdf>

The main advantages of VMware guest backup is that it allows for support of partial or file-level restores across different operating systems and lets users leverage identical backup methods for physical and virtual machines. There is no requirement for advanced scripting or VMware software knowledge — this means unchanged day-to-day procedures for backup. The application consistency that is typically a requirement for production applications is delivered through this method.

Guest-level backups have been qualified in conjunction with VMware Distributed Resource Scheduler, vMotion, and High Availability solutions that allocate virtual machines based on demand or move virtual machines for business continuity and disaster recovery. As virtual servers move, NetWorker is still able to backup and recover as needed.



- Individual backup for each virtual machine
 - VMs treated like physical machines
 - Software installed on guests
 - Backup sent through LAN
- Simple integration and familiar workflows
- Optimal protection and recovery for applications
- VMware Distributed Resource Scheduler, vMotion, and High Availability Support
- Deduplication relieves bottlenecks

Figure 13: NetWorker - Guest-Level Backup

EMC NetWorker Module for Databases and Applications

EMC NetWorker Module for Databases and Applications integrates with Data Domain through DD Boost within the NetWorker Module for Databases and Applications 1.2 (and higher) and the NetWorker Module for SAP 4.2 (and higher). NMDA provides increased performance, seamless replication management, simplified administration, and it enables application-consistent backups by communicating with the applications running in the virtual machine without bottlenecking ESX resources.

- **Integrates with:**
 - Oracle Recovery Manager (RMAN)
 - DB2 backup and restore API
 - Informix XBSA
 - Sybase Backup Server API (BSA)
 - Lotus C API
 - Simplified, common configuration steps and workflows
- **Wizard-driven configuration**
- **Reduces backup impact and required storage**
 - Integrated deduplication
 - Snapshot management
- **Delivers value-added features**
 - Event-based backup
- **Integrates with VMware**



Figure 14: NetWorker Module for Databases and Applications

Long-Term Retention of Backups

Many companies have minimized the use of tape by deploying deduplication storage for backup and operational recovery –Data Domain systems have been the market leaders in this category. In general, operational recovery includes retention periods from a few weeks to a few months. For longer retention requirements, most users simply keep tape backups longer, while few use a specialized disk-based archive approach. Long term retention of data is increasingly important in VMware environments because of compliance requirements.

While tapes offer portability and low acquisition cost, there are several drawbacks to using them for data retention:

- The perceived capital cost advantages of tape cartridges are often outweighed by the costs of tape management and operating expenses
- Recovering from tape is unreliable and time consuming especially if the tape is stored offsite
- Evolving tape drive technology can require expensive and time-consuming bulk migration of old tapes every 5-7 years

An alternative to using tape and backup processes for long-term retention is to use specialized archiving processes or applications to move data from primary storage and its corresponding backup process to an archive platform.

EMC Data Domain Archiver

EMC Data Domain Archiver (DD Archiver) is the industry's first system for long-term retention of backup and archive data, supporting today's use of backup processes for data retention and evolving adoption of archiving workloads. DD Archiver extends a mid-range Data Domain system configuration, which is already widely deployed and proven in backup environments and transparently incorporates a large-capacity tier dedicated to static data. New architectural enhancements allow the system to incorporate very large capacities that can expand over time, reducing system cost, and ensuring long-term availability and integrity of data. When archiving software approaches do fit the environment, users can target their archive workloads to the same DD Archiver in the same infrastructure.




Figure 15: Data Domain Archiver

With the passing of time and enforcement of retention policies, the system can expand by incorporating additional archive units. In this fashion, data gets laid out into archive units in age order.

The architectural innovation is to separate short-term data from aging long-term data into different tiers, and archive data gets laid out into archive units in age order. The active tier holds the short-term data for disaster recovery purposes, while the archive tier holds static long-term retention data. This separation of data types is what allows the impressive scalability of the archive tier, while still keeping the performance required for the active tier.

Centralizing Management of Backups

Centralized management and reporting provide the operational control and transparency needed for today's compliance challenges and assures predictable data recovery. EMC backup and recovery solutions provide powerful centralized management that delivers the needed oversight for solid data protection.

EMC Avamar

Avamar is tightly integrated to Data Domain systems for increased scalability and through its management console it manages all scheduling, policies, retention and recoveries for both the Avamar Data Store and the Data Domain systems.




Figure 16: Avamar Centralized Management

EMC NetWorker

Today's data protection solutions need to operate across a very diverse landscape—managing multiple tiers of data protection from tape, to disk, to snaps, across a variety of applications within both physical and virtual environments, and providing comprehensive scheduling and policies for both operational and disaster recovery scenarios. To drive complexity and cost out of this environment, you need common management.

NetWorker is a powerful application that provides centralized control and command (a single management interface) necessary to provide greater simplification and seamless management to meet today's new requirements.




Figure 17: NetWorker - Unified Data Protection

Not only does NetWorker unify existing infrastructure that has evolved over time, but it helps to bridge the transition from traditional to next-generation backup methodologies.

EMC Data Protection Advisor – Data Protection Management Software

EMC Data Protection Advisor (DPA) provides a single, comprehensive view of the entire backup infrastructure through automated data collection, analysis, alerting and reporting that leverages this data for key backup management functions.

Data Protection Advisor transforms volumes of disparate data across locations and backup solutions into actionable business information that allows companies to:

- Lower costs through the improved use of their infrastructure, avoiding unnecessary purchases and reducing manual effort
- Improve compliance and lower risk through better visibility and assurance that critical data is protected
- Reduce complexity with a single console to provide an integrated, automated view

Users typically see payback from using Data Protection Advisor in 12 months or less—with ongoing total cost of ownership benefits.

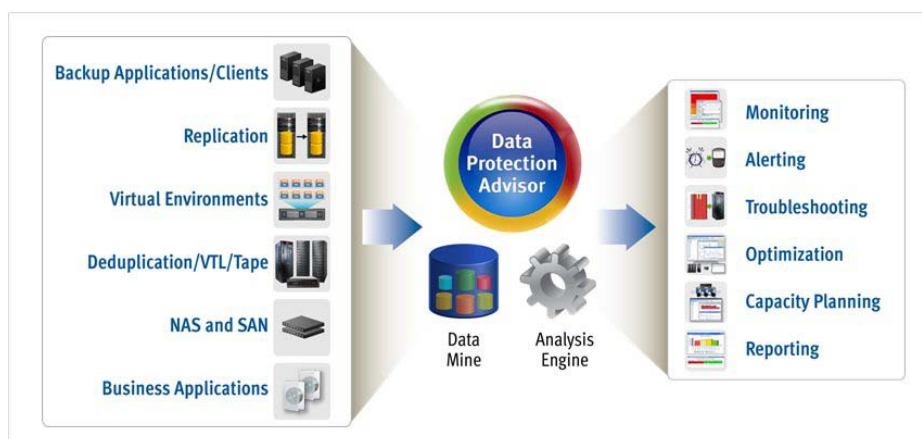


Figure 18: DPA – Gain Complete Visibility of your Data Protection Environment



Figure 19: EMC Data Protection Advisor

- **Increase visibility**
Monitor, analyze and report across physical and virtual backup environments with a single, unified product
- **Reduce costs**
Consolidate information enterprise wide to reduce administrative efforts, accelerate problem resolution and forecast resource capacity planning
- **Prove protection**
Understand potential risks and data protection exposures across multiple backup products
- **Lower complexity**
Use an automated, unified approach to managing the data protection environment, greatly simplifying the monitoring process and tuning the environment to maximize utilization

Conclusion

EMC integrated backup and recovery solutions provide a comprehensive approach to data protection for challenging VMware environments. Avamar and NetWorker, with VADP, provide efficient off-host backup that is integrated with vCenter to protect VMDK images for disaster recovery. Avamar and NetWorker likewise protect applications with intelligent agents that can be optionally configured with Data Domain deduplication storage systems. By deduplicating the data at the source with Avamar or at the NetWorker client with the DD Boost integration, application consistency is preserved without overwhelming shared VMware resources.

The centralized management capabilities of Avamar and NetWorker, in conjunction with the cross-domain monitoring, analysis and reporting features of Data Protection Advisor ensures efficient operational control of all data protection components. It is for all of these reasons that more people choose to build their VMware backup solutions using EMC products and technology.

Appendix

Backup Application Software Option Settings Recommended for use with Data Domain systems

Most commercial backup application software has options that can interfere with the optimal deduplication ratios that are attainable using Data Domain systems. Table 2 provides recommendations for these settings to provide the best deduplication ratios and backup performance.

Table 2: Data Domain – Best Practice for Backup Application Software Settings

Option	Setting	Explanation / Benefit
Compression	Disabled	<ul style="list-style-type: none">• Backup software compression should be disabled because the Data Domain system can fingerprint unique data segments more efficiently for deduplication if the data segments sampled are not already compressed• Backup windows can be extended and cpu performance can be impacted on the backup client if the backup software is tasked with performing compression• Local compression is provided for on the Data Domain storage system
Encryption	Disabled	<ul style="list-style-type: none">• Encrypted files are by definition, unique. The encryption software that is part of the backup application will create unique files, on-the-fly for each backup, defeating the deduplication capabilities of the deduplication storage system.• Optional software is available for Data Domain systems to provide for encryption at rest
Multiplexing	Disabled	<ul style="list-style-type: none">• Multiplexing of backup images was designed to provide better ingest rates for physical tape drives that need to be kept streaming to keep the tape from stopping. When physical tape motion is halted, the tape must reverse in order to register the tape's position to resume writing. This behavior is known as, "shoe-shinning" (or "back-hitching").• Disk drives do not have shoe-shinning as an issue and multiplexing can interfere with the fingerprinting process and deduplication ratios
Deduplication	Disabled	<ul style="list-style-type: none">• Disabling deduplication from the backup application software will provide better performance and allow the Data Domain system to offload this work• Data Domain systems are optimized to provide the very best ingest performance and deduplication ratios

Additional Resources

EMC Avamar Product Page: <http://www.emc.com/backup-and-recovery/avamar/avamar.htm>

EMC NetWorker Product Page: <http://www.emc.com/backup-and-recovery/networker/networker.htm>

EMC Data Domain Product Page: <http://www.emc.com/backup-and-recovery/data-domain/data-domain.htm>

EMC Data Domain Boost Data Sheet: <http://www.emc.com/collateral/software/data-sheet/h7034-datadomain-boost-sw-ds.pdf>

EMC Data Protection Advisor Product Page: <http://www.emc.com/backup-and-recovery/data-protection-advisor/data-protection-advisor.htm>

IDC Study – Worldwide Purpose Built Backup Appliances:
http://www.emc.com/collateral/analyst-reports/11530-idc-ww-pbba-2011-2015-forecast.pdf?cmp=nlc-dedupenews-august-idcshowswhy&AM_ID=4998589&ActID=27626&M=89c5ccf3-143f-4272-ae99-d61e4305e999

EMC Data Domain Data Invulnerability Architecture: Enhancing Data Integrity and Recoverability: <http://www.emc.com/collateral/software/white-papers/h7219-data-domain-data-invul-arch-wp.pdf>

VCB End of General Support: <http://www.vmware.com/support/policies/lifecycle/vi/eos.html>

VMware vStorage VMFS: <http://www.vmware.com/products/vmfs/overview.html>