

EMC Celerra NFS for VMware

Deploying virtualized business-critical applications requires infrastructure-quality NFS

The Big Picture

- No-compromise infrastructure for VMware deployments, designed to deliver five-nines (99.999 percent) availability through advanced hardware and software architecture
- Industry-leading scalability with up to eight X-Blades and 900 TB of storage, supporting hyper-consolidation of more virtual servers and applications onto a single storage platform
- Advanced functionality at no additional cost such as writeable snapshots, file system deduplication, automated volume management, and virtual provisioning in an easy-to-use, easy-to-install, NAS storage offering
- VMware vCenter integrated management for storage provisioning, optimization, and virtual machine (VM) cloning
- Industry-leading support for VMware Site Recovery Manager through SRAs and vCenter failback plug-in
- Flexible connectivity with unified storage: NAS (CIFS and NFS), iSCSI, and Fibre Channel
- Optimized performance-sensitive portions of VMs and virtualized applications

Gaining the flexibility of NFS without compromise

EMC® Celerra® was originally developed to provide customers with the same levels of availability, reliability, performance, and scalability for file shares that they had come to rely on in their EMC Symmetrix® block-storage systems. Specifically, general purpose and appliance servers were not sufficient for running mission-critical engineer and manufacturing operations. Celerra has since evolved to be the industry-leading¹ NAS system. Moreover, the attributes that make Celerra the preferred infrastructure for mission-critical file-sharing applications also make it the best storage infrastructure for deploying VMware®.

Infrastructure for VMware

Key attributes of an NFS infrastructure for VMware include:

- **Availability:** Storage systems must deliver high availability and advanced features to ensure virtualized application performance.
- **Protection:** The NFS infrastructure must maintain virtualized application availability and protect information without adversely affecting business operations.
- **Efficiency:** One of the key benefits of VMware is improved provisioning and operational efficiency. The storage infrastructure must be able to extend efficiencies to storage management.
- **Flexibility:** Virtualized infrastructures enable rapid changes in application deployments and resource allocation. The NFS infrastructure must be able to respond to changes in storage capacity, performance, and protection requirements.
- **Manageability:** One of the benefits of deploying VMware on NFS is consolidated storage management. The NFS solution must provide simplified and automated management and extend this by integrating control into the VMware vCenter™ Console.
- **Service and support:** The infrastructure provider must be capable of ensuring that customers realize the value of their virtualized infrastructures while reducing risk and accelerating projects.

Availability


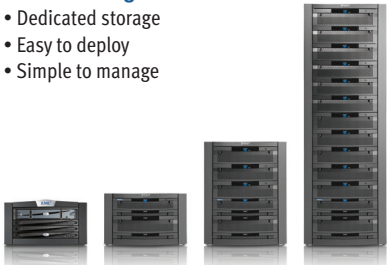

Celerra provides measured five-nines (99.999 percent) data availability across thousands of customer installations. This proven performance is required when deploying virtualized infrastructures where a service disruption can bring down hundreds of application or thousands of virtual desktops.

The modular architecture is composed of file-serving blades (X-Blades) and storage infrastructure and provides redundancy and scalability to hundreds of terabytes of NAS storage and thousands of disk drives. Celerra delivers the flexibility to increase capacity, performance, or both to meet the dynamic needs of a virtualized data center.

The Celerra DART operating system delivers data availability and protection as top priorities. The file system is structured to minimize file fragmentation. This reduces file system maintenance, provides predictable read/write performance, and enables over 90 percent storage utilization, optimizing the storage investment and minimizing the number of file systems managed in the VMware environment.

Further, the file system uses metadata logging to ensure data integrity and minimize restart time after a service disruption. You can build large file systems without worrying about a momentary power disruption that could result in hours of data unavailability while performing file system recovery.

Figure 1:
Celerra Family of Unified and Gateway Systems

	Unified Storage <ul style="list-style-type: none"> • Dedicated storage • Easy to deploy • Simple to manage 				Gateway <ul style="list-style-type: none"> • Shared storage • Add NAS, iSCSI to Fibre Channel SAN • Highest storage usage • Single management 	
						
	NX 4	NS-120	NS-480	NS-960	NS-G2	NS-G8
Number of X-Blades	1 or 2	1 or 2	2 or 4	2–8	1 or 2	2–8
Storage	Integrated	Integrated	Integrated	Integrated	FC SAN	FC SAN
Maximum file capacity (X-Blade/system)	16 TB/ 32 TB	16 TB/ 64 TB	64 TB/ 192 TB	128 TB/ 896 TB	64 TB/ 128 TB	128 TB/ 896 TB

Celerra is available as a gateway system, utilizing EMC CLARiiON® and Symmetrix storage via 4/8 Gb FCSW SAN infrastructure or as an integrated unit, packaged with the latest CLARiiON technology to provide the industry-leading unified storage platform.

Protection

Celerra also provides industry-leading technologies to protect information from loss due to application or user error and site disasters. Important data protection attributes include:

- Managing the protection process with minimal impact to production environments
- Quick, efficient, and flexible recovery
- Infrastructure that minimizes capital and operational expenses

Working with other EMC products integrated with VMware, Celerra meets and exceeds these requirements to provide the most comprehensive infrastructure, data, and site-protection solutions. Using Celerra, the following capabilities are easily deployed to meet varying protection requirements:

EMC Celerra SnapSure™ software creates copies of file systems within the Celerra system, which can be used for backups and quick recovery of deleted files, virtual machines (VMs), or file systems. These copies are also available for testing or other activities where copies of “production” data will enhance business operations. When deploying VMware on NFS, SnapSure makes it easy to make copies of all or a part of the entire VMware infrastructure.

SnapSure uses a pointer-based implementation, so the space consumed by copies is directly proportional to the change rate of the file system. SnapSure uses a Copy on First Modify (COFM) technology. As data is written to the file system, data that would have been over-written is moved to a private storage area.

COFM offers several benefits, including:

- Only production data is stored in the product file system, so the production file system does not fill up with “copies of old data,” which causes unpredictable file system growth and affects file system operations.
- Moving old data out of the production file system minimizes fragmentation thus eliminating production read performance impacts.
- Deleting old copies of file systems does not impact the production file system.

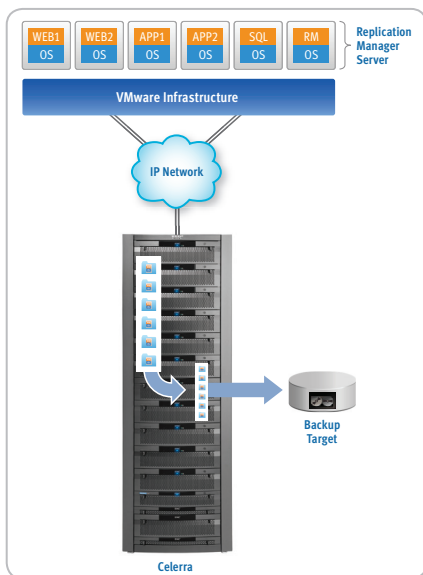


Figure 2:
Replication Manager with VMware on NFS

SnapSure is included in the base Celerra software package and includes a GUI-based scheduler plus CLI to integrate with backup or other external data operations. A SnapSure copy of a VMware data store will produce “crash consistent” copies of the VMs. Recovery from these copies will be similar as recovering after a hard “crash” of the system. This is adequate for many backup/restore requirements. However, some applications—such as databases, messaging, and transactional programs—gain significant management benefits from “application-consistent copies.”

EMC Replication Manager provides host-based management of array-based copies to produce consistent copies of application and VM data. Replication Manager is integrated to work with VMware deployed on Celerra NFS. A single Replication Manager console can be used to manage VMware infrastructures deployed on NFS, Fibre Channel, and/or iSCSI on EMC Symmetrix, CLARiiON, and Celerra.

Replication Manager is integrated with application APIs such as Microsoft® VSS, VDI, and Oracle RMAN to produce application-consistent copies without data or business operations.

Replication Manager is also integrated with the VMware vStorage APIs, using vmsnap to produce application-consistent copies of VMs; and in the case of Microsoft Windows VMs, enabling consistent copies for applications like SQL, Exchange, and SharePoint with Windows VMs.

The copies made with Replication Manager can be a source for backup processing via EMC Avamar®, EMC NetWorker®, or other backup software, as well as business repurposing for test or reporting operations.

EMC Celerra Replicator™ replicates one or more file systems to another Celerra system to protect information from a site failure. Replication is set up per file system, based on the business recovery-point-objectives (RPOs). For instance, critical file system(s) may be set up with an RPO of 10 minutes, i.e., no more than 10 minutes of information will be lost due to a catastrophic site failure. Other file systems with less important information will be set up with RPOs of hours or perhaps a full day for information that simply needs to be “backed up” to the disaster recovery site. Celerra Replicator schedules replication sessions based on the defined RPO requirements, data change rates, and available IP network bandwidth to maintain data availability at the disaster recovery site.

With the addition of VMware Site Recovery Manager (SRM) support for NFS in VMware vSphere™ 4, Celerra Replicator now supports SRM in NFS environments. A storage replication adapter (SRA) for Celerra Replicator is available at www.VMware.com. The SRA extends EMC’s industry-leading SRAs produced for other replication technologies, such as EMC SRDF®, EMC MirrorView™, and RecoverPoint. The SRA enables SRM to leverage Celerra Replicator to execute VM restart at the disaster recovery site. In addition, SRM can use SnapSure on the target Celerra to enable disaster recovery testing without affecting production operations or data replication.

In addition to the SRA, a vCenter plug-in is available to automate many of the tasks required to fail back. When operations need to be returned to the primary location after a “disaster event,” the vCenter Failback plug-in will: 1) resynchronize the production data storage with the disaster recovery data storage, 2) enable the storage at the production site so VMs can be restarted, and 3) resume data replication from the production facility to the disaster recovery facility. Automating the restart of the storage operations at the production facility greatly simplifies the processes required for the VMware administrator to “go home” successfully.

Efficiency

One of the key benefits of VMware is resource efficiency and optimization. The VMware ESX® hypervisor enables you to optimize computer resources by consolidating application workloads. VMware tools like Distributed Resource Scheduler (DRS) and Dynamic Power Management (DPM) enable the system to adjust resource allocations automatically to changing requirements. The storage infrastructure needs to provide similar capabilities and benefits.

Provisioning storage for a VMware environment is a balance between optimizing storage capacity to minimize cost and providing adequate performance for consolidated and dynamic workloads. Storage capacity efficiency can best be achieved by allocating only what is needed and minimizing growth as evidenced in the following:

Celerra virtual provisioning is an advanced implementation of thin provisioning. With virtual provisioning, only the space in the file system that is consumed by VM data is allocated from the storage pool.

Consider a scenario where a VMware administrator is starting a new project. When the project is completed, it is expected to require about one TB of file system space. The administrator creates a one TB virtually provisioned file system and creates three VMs. Assuming each is about 20 GB, the total space consumed will be about 60 GB, providing space for the three VMs plus a little extra for growth. As the administrator adds more VMs, more space is consumed. If the project never reaches the one TB expectation, the extra storage space is available for other projects.

Celerra file system management provides advanced reporting on data consumption rates, including predictive reporting to estimate when space is expected to be exhausted. It will also notify the administrator when watermarks have been exceeded to provide adequate time to add storage and determine why growth rates have exceeded expectations.

Celerra data deduplication leverages efficient hashing and compression algorithms to reduce the physical storage used by files in the file system. Celerra data deduplication is optimized to work with VMDK files in an NFS data store, reducing space consumption by up to 50 percent. The overhead of accessing the compressed VMDK is less than 10 percent, which is insignificant for most VMs. However, unlike other VM storage deduplication solutions, Celerra data deduplication enables the VMware administrator to select which VMs in the data store to compress. If the overhead is too much for any VM, the administrator can simply uncompress it on the fly.

Fast virtual machine cloning enables the VMware administrator to make copies of VMs that consume the minimal amount of space required to present the unique image to ESX via NFS. Fast cloning is similar to VMware Linked Clones, except Celerra performs all processing and does not use ESX resources. The administrator can create a single or hundreds of copies of a VM—for instance when deploying VMware View™. From the Celerra perspective, these clones are simply files in the data store and are managed like any other file.

The combination of virtual provisioning, deduplication, and thin cloning makes it easy for the administrator to optimize storage provisioning and consumption.

Performance optimization is the other key aspect of VMware storage management. Provisioning storage that delivers service levels that meet performance requirements is critical for maximizing the storage investment's value while reaping application consolidation benefits. VMware deployed on Celerra NFS fully exploits the value of industry-leading storage capabilities.

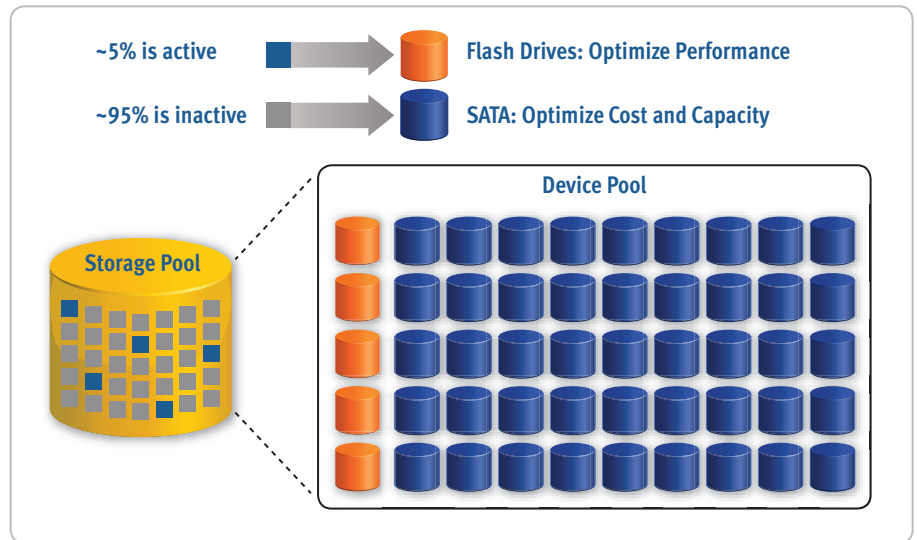
Storage tiering enables the administrator to place data on drive technology that matches the performance requirements. Celerra supports SATA, Fibre Channel, and Flash drives within the same array. In addition, Celerra supports varied RAID types, including 1/0, 5, and 6 with multiple RAID group geometries. The flexibility of having multiple drive and RAID types available means the administrator can match storage implementation to performance requirements.

Fully automated storage tiering (FAST) is a revolutionary new technology that automates performance optimization. While traditional storage tiering enables users to put an entire LUN (logical disk) on a storage tier, sub-LUN FAST, which will be available later in 2010, will move much smaller pieces of storage around in the array to respond to changing I/O requirements. This increased granularity will greatly improve the efficiency of storage tiering.

A large portion of a VMDK file that stores a Windows image tends to be inactive, while other portions can be extremely active. FAST Sub-LUN tiering places the very active portions of the VMs on Flash drives, leaving the bulk of the storage on SATA or FC drives to optimize the storage investment. Some examples of the data that will be accessed via Flash drives with Sub-LUN FAST are:

- VMware cache page files
- Windows swap/page files
- Database indexes/active tablespaces
- Other active data that would benefit from decreased response times

Figure 3:
FAST Performance Optimization



Sub-LUN FAST is fully automatic, so the storage layout process needs to address what percentage of the Virtual Storage Pool should be SATA, FC, or Flash.

The combination of Celerra and EMC storage array functionality makes maximizing storage efficiency easy and automated.

Flexibility

A storage solution for VMware must be flexible enough to easily accommodate changes in infrastructure requirements. NFS will work as a storage protocol for most VMware deployments, but there are use cases when requirements will call for deploying iSCSI, FC, or FCOE. Celerra, in combination with storage array technologies, makes it simple to implement a multi-protocol storage infrastructure.

Tiered Storage Protocol Support. Celerra gateway systems use FCSW to attach to EMC CLARiiON or Symmetrix storage arrays. These arrays support iSCSI, FC, or FCOE connectivity. Celerra unified storage systems are packaged with the latest CLARiiON array technology. Additional connectivity can be enabled on the CLARiiON system to provide the industry-leading, mid-tier, block-storage protocol support.

All of the features on the latest CLARiiON technology to support VMware on block protocols are available including:

- Virtualization Aware Management with EMC Navisphere®
- Automated storage path management with EMC PowerPath®/VE
- Replication with RecoverPoint including VMware SRM integration
- StorageViewer vCenter plug-in
- Replication Manager for VMware
- Fully automated storage tiering (FAST)

Performance and capacity scalability are critical to meet the requirements of virtualized environments. Virtualization enables IT to respond to new technical and business opportunities much quicker than in non-virtualized data center environments. It is critical that the storage infrastructure is able to scale to take on new projects without affecting existing applications or impeding virtualization progress. Celerra has modular components based on a common architecture, so the environment can be expanded by adding hot pluggable components and data-in-place upgrades. For example, a project could start with an NS-120 system with one active X-Blade with a few TBs of storage and grow to an NS-960 with seven active X-Blades and 900 TBs of storage.

The NS-480, NS-960, and NS-G8 models offer multiple X-Blades configurations. Each has a minimum configuration of two X-Blades. X-Blades can be added to increase performance, capacity, or both. For instance, if it is determined that there is adequate storage capacity, but additional NFS performance is required, an X-Blade could be added to the system to increase performance. The ability to scale

capacity and performance independently increases the flexibility of the infrastructure and avoids expenditures for unnecessary components resulting from rigid architectures.

Many VMware infrastructures deploy Celerra with more than two X-Blades in order to achieve workload isolation. An NS-480 with three X-Blades, for example, could use one blade for VMware NFS infrastructure, another for CIFS consolidation for user or windows share data, and a third for test/development. With Celerra, an NFS data store can be moved from a test/development environment to production by simply moving the access point to the storage—no physical data movement is required.

The capacity and performance scalability of the Celerra systems provide a common infrastructure for deploying VMware on NFS, iSCSI, FC, or FCoE that will not inhibit the growth of virtualization projects or fall short of performance requirements.

Manageability

Celerra offers significant functionality for deploying VMware where ease of use is a key attribute. Celerra Manager offers a comprehensive, full-function GUI, including at-a-glance reporting, wizards, task schedulers, and call-home/remote diagnostics. Celerra Manager provides the storage administrator with the tools to make intelligent decisions, be pro-active, and execute storage management tasks simply and easily. Providing the VMware administrator similar capabilities to manage storage infrastructure components is important for efficient virtualized infrastructure management.

EMC Celerra Plug-in for VMware enables VMware administrators to manage NFS data storage in the Celerra from the VMware vCenter 4 console. The “NFS plug-in” provides three key functions:

Provisioning: Create an NFS file system and automatically mount it to the ESX systems in the cluster. File systems created with the NFS plug-in will be automatically deployed with EMC and VMware best practices, including Celerra Virtual Provisioning. File systems can also be expanded without disruption to the VMware environment.

Virtual machine cloning: In addition to management of the Fast Clones (as described earlier), the NFS plug-in supports the creation of full clones. This can be used to create a copy of a VM to another file system, such as creating a VM from a gold copy in a protected file system.

Virtual machine compression: Leveraging Celerra deduplication, the plug-in manages the reduction of space consumed by active VMs. The VMware administrator can compress a single VM, all VMs on a server, all VMs in a folder, or all VMs in the ESX cluster. The administrator can also decompress any VM on demand.

In addition to the EMC Celerra Plug-in for VMware, a VMware SRM failback plug-in (mentioned earlier in this document) is available to manage the storage tasks of returning to the production facility after a disaster recovery failover scenario.

Maximize the benefits of Celerra with EMC Global Services

EMC delivers the full complement of services for Celerra products to ensure they perform as expected in your storage environment while minimizing risk to your business and your budget. Expert planning, design, and implementation services help you quickly realize the value of your investment in your environment—no matter how simple or complex.

The Celerra QuickStart services rapidly deliver a fully functioning NX4, NS-120, NS-480, or NS-960 unified storage system, serving files to authenticated users in a production environment. The service includes configuration of network interfaces, file systems, and other software; implementation and testing of your hardware; and concludes with a functional overview of product features delivered to your staff. EMC can assess your environment to determine the optimum configuration and integration of your Celerra storage platform into your infrastructure, create a detailed technical design, and review the lifecycle of your data assets to help you define and develop the ideal IT organization for your storage environment and best-practice policies to support it.

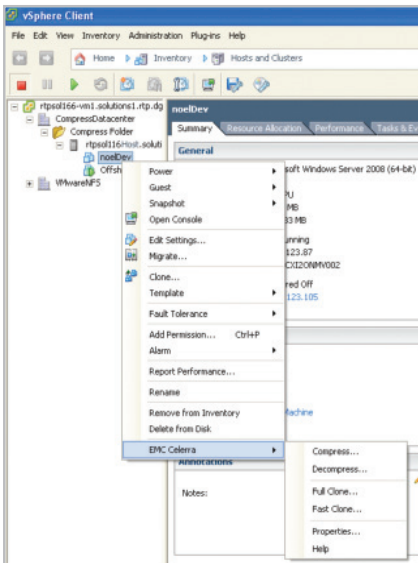


Figure 4:
EMC Celerra Plug-in for VMware

EMC Consulting helps you unleash the value of virtualization with an integrated approach across server, storage, application, network, and desktop assets. Our infrastructure, application, and business specialists accelerate the adoption and integration of operation best practices and processes for the virtual infrastructure, encompassing backup, disaster recovery, and integrated management of physical and virtual assets. EMC Consulting services for VMware environments include virtualization assessments, strategy development, operation readiness, enterprise data center infrastructure and operations planning and design, VMware View (VDI) planning, business continuity for the virtual infrastructure, program management, and virtualizing enterprise applications such as Microsoft Exchange and next-generation data center.

EMC Proven™ Solutions leverage key virtualization features within EMC and non-EMC products, enabling you to design and build out the virtualized data center with predictable results. EMC extensively tests combinations of products and develops best practices and reference architectures to minimize your risk of development while accelerating ROI. EMC Proven Solutions span a wide range of business objectives such as virtualizing tier 1 applications—including Microsoft Exchange/SharePoint/SQL, Oracle, and SAP—and addressing your specific business needs including performance, backup/recovery/archive, and business continuity/disaster recovery for mission-critical applications, unified/tiered storage, management, and security.

EMC Customer Service—six-time winner of the SSPA STAR Award for outstanding mission-critical support—helps you keep your information available 24x7 to deliver competitive advantage and drive revenue. EMC Education Services drives the value of your investment with a comprehensive portfolio of customer courses.

Ask your EMC sales representative about our full spectrum of services that can benefit your organization.



EMC Corporation
Hopkinton
Massachusetts
01748-9103
1-508-435-1000
In North America 1-866-464-7381
www.EMC.com

Take the next step

For more information on how EMC Celerra unified storage platforms can meet your networked information-sharing needs and bring increased value to your business, contact your EMC sales representative or authorized EMC value-added systems integrator. Or, visit our website at www.EMC.com.