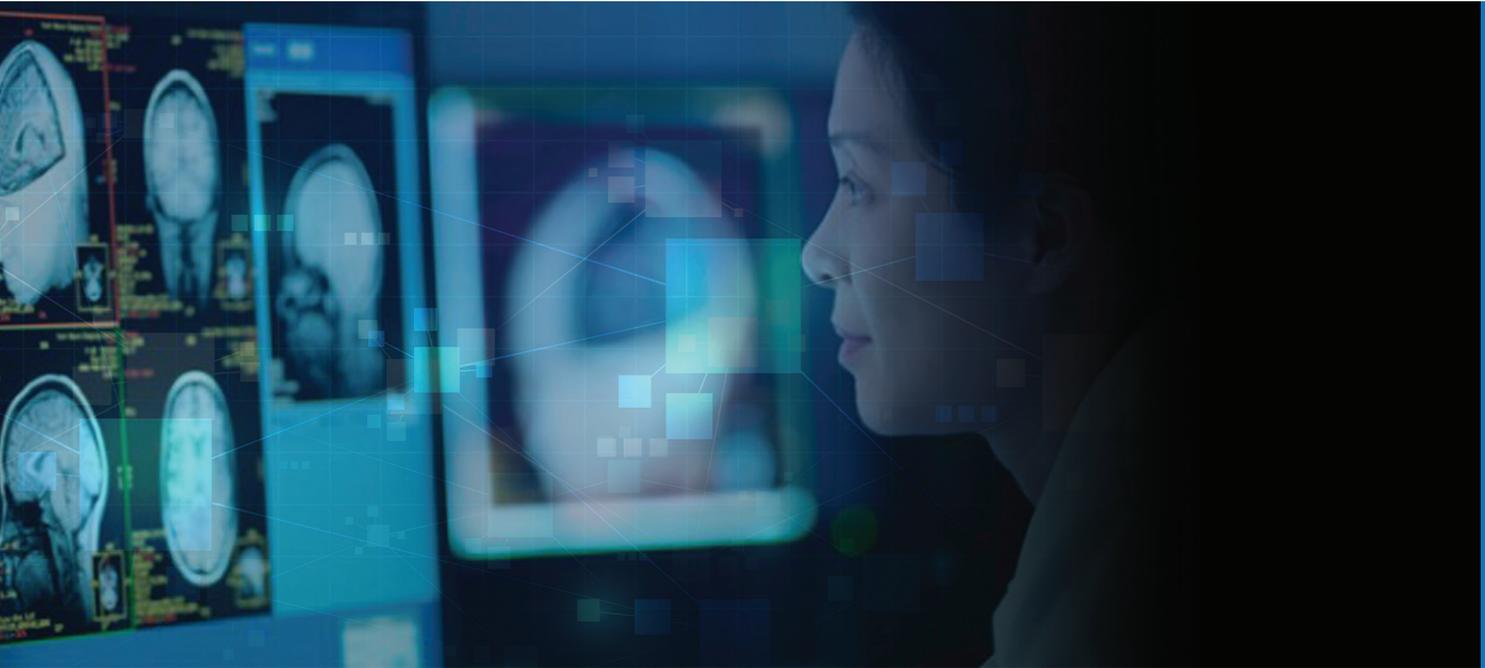


Hyperconvergence and Medical Imaging



SERVERS

DATA STORAGE

PCs & WORKSTATIONS

 **nfina**

HYPERCONVERGED STORAGE FOR HEALTHCARE



Surveys show the top three challenges facing medical IT departments with regard to data storage are:

1. Lack of Flexibility

Requirements for dedicated hardware and difficulty integrating and sharing the data with entities utilizing different environments make existing systems cumbersome.

2. Poor Utilization

Existing and traditional systems for file storage have capacity requirements that either push the limits of fault tolerance, or simply make fault tolerance impossible. Also, with traditional storage strategies, over provisioning is a necessity, which translates to considerable expenditures on unused hardware.

3. Inability to Scale

RAID based arrays, as well as other hardware solutions are difficult and expensive to expand, and redundancy requirements result in an inefficient use of storage space.

HEALTHCARE DATA

According to one report, in 2014 the volume of healthcare data was 153 Exabytes. It is projected to explode to 2,400 Exabytes by 2020. Many hospitals, radiology clinics, and other entities that produce large amounts of medical imaging are scrambling to find efficient and cost effective storage strategies to cope with this massive increase in data.

2014 – 153 Exabytes

2020 – 2400 Exabytes

In the recent history of medical imaging, PACS (Picture Archiving and Communication System) and DICOM (Digital Imaging and Communications in Medicine) are technologies that have been developed to manage the growing amount of medical image data, and both technologies made sharing and storage of images an easier task.

These technologies were improvements, but problems still existed. Different vendors created their own implementations, and many of these implementations used proprietary hardware and software and were not open systems. This locked users into a particular vendor's hardware. Given the rapid pace of changing technology, being locked into a hardware solution is a real problem. It is a fact of life that something newer, faster, and better will debut in the very near future. Consider a state-of-the-art purchase just a few years ago – it would have included tape drives, optical jukeboxes, and other hardware that were made obsolete by advances in technology in a very short time frame. Rapid and continuing advances in storage technology means that adopting newer, faster, and more reliable hardware is not a nicety, but rather a certainty. Being locked into specific hardware means it will be difficult, if not impossible to take advantage of advances in technology without abandoning the existing system and opting for the dreaded “forklift upgrade”.



Enter the Vendor Neutral Archive (VNA). The definition of a VNA is “a medical device that stores medical images in a standard format with a standard interface, such that they can be accessed in a vendor-neutral manner by other systems”. A true VNA will eliminate a number of barriers for medical imaging and medical IT professionals. It will remove barriers to sharing information between different vendor implementations of medical data creation, sharing and storage. It means that facilities using different vendors’ hardware can still share medical information, and it means that as storage technologies evolve, medical IT personnel can easily update storage technology and scale storage capacity while considering only performance and value – not the constraints of proprietary technologies. Medical IT professionals must be diligent in researching providers, because some vendors who claim to have VNA solutions still use a proprietary archive, which means they are not really vendor neutral at all. Taking the steps to assure that a system under consideration is truly 100% vendor neutral means that a medical IT professional’s job will be much easier when it comes time to upgrade or add capacity.

Virtualization

Virtualization removes the need to deploy specific instances of hardware and software, instead creating virtual machines within servers. The virtual machine can be equipped with the software and resources (applications, memory, processing power, etc.) necessary for the intended task, and when no longer needed, the virtual machine can be deleted and the resources it consumed are returned to the resource pool in the server. This is incredibly more efficient and versatile than purchasing and deploying individual machines and installing software on them. Virtualization enables a doctor or technician to move from location to location within an enterprise, log on to a thin client, and have the same resources, GUI, data, and other tools necessary without having to carry a laptop or other hardware. It is difficult to make a case against virtualization on the compute side of the equation.

Storage as a Choice: Hyperconverged vs. Storage as an Appliance

With regards to storage, there are choices. If there is one hard and fast rule for evaluating IT solutions of any type it is this: Choose the system, hardware, and software for the job, for the specific application, and for the specific organization. This white paper will look at two storage strategies, hyperconverged storage and storage as an appliance.

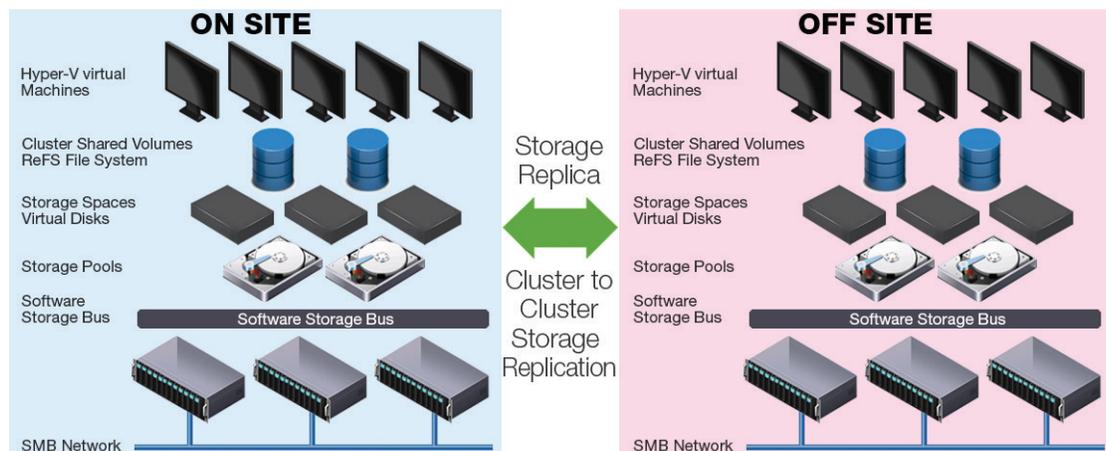
Hyperconvergence integrates compute, storage, and virtualization resources, and manages them with a software appliance. Hyperconvergence enables IT professionals to do considerably more while using fewer resources. Less hardware is required, and by simplifying software and hardware management, less of the IT department’s valuable time is required. It provides more availability, more resiliency, and better disaster recovery, all with less IT labor and hardware expenditures. When more hardware resources are necessary, new nodes are simply added to the cluster. It is important to note that when a node is added, it must be identical to existing nodes. Technology advances rapidly, and often within a short period of time (1-3 years) installed nodes may be obsolete, making the addition of an identical node difficult or impossible. Once the initial deployment is completed, operations can be handled remotely. When a new node is required, all that is necessary is the installation of the hardware, cabling, and power – configuration can be done by remote IT personnel.

Hyperconverged systems balance the computing and storage load among the installed nodes. A side benefit of this balancing act is increased availability. If a node fails, the load is redistributed among the functioning nodes. This has important differences from traditional methods of fault tolerance. Before hyperconvergence, for a server to be fault tolerant, it generally required a second server mirroring the operations of the first. If the first failed, the second server would take over. With hyperconvergence, no duplication of hardware or software is required. When a failure occurs, the failed machine's load is redistributed to the functioning machines.

Hyperconverged storage is relatively resilient and reliable. Storage is distributed among the nodes. In a traditional system with a NAS or SAN using RAID, disk failure necessitates that IT personnel replace the disk as quickly as possible, because the system was not only vulnerable, but often was operating in a degraded mode as a result of the disk failure. Some RAID levels can tolerate two disk failures, but an additional disk failure means that data will be lost. With Hyperconverged systems, a failed disk is not such an urgent situation; the hypervisor in charge of storage simply redistributes storage to other functioning resources.

Hyperconverged systems are readily scalable. In a hyperconverged system, the hardware needs only to be installed and cabled, and the minimal configuration (if any is required) can be done quickly and from a remote location.

Hyperconverged systems are relatively efficient. Rather than deploying a dedicated server supporting a department or function, a server that often is over provisioned and under utilized, a hyperconverged approach takes existing resources and utilizes them more fully. Instead of deploying multiple instances of individual servers and storage supporting individual departments or applications, a hyperconverged approach employs a cluster of servers to support the enterprise, and they are usually fewer and more completely utilized than in traditional systems. But, hyperconverged systems utilize the compute power of the server for storage. With dedicated storage appliances, storage computation is off-loaded to the appliance, freeing the server to devote more computing power to applications. Hyperconverged systems usually provide an increase in performance at the end user level as well. Some medical centers have observed as much as a 30% improvement, resulting in a time savings of up to 45 minutes a day – a considerable productivity boost for busy physicians, technicians, and other consumers of a radiology center's data.



Hyperconverged network for a radiology clinic utilizing Microsoft® Storage Spaces® and Hyper-V®

Storage as an Appliance

Hyperconverged systems lend themselves to small to medium sized organizations. For large enterprises keeping compute, networking, and storage functions apart is the better, if not only choice. Depending on the enterprise and number and types of applications, the migration of existing applications may prove costly and difficult, even to the point of outweighing the eventual benefits of hyperconvergence.

A storage appliance provides and manages data for other network attached computing devices. Appliance types include NAS and SAN devices. They are typically added where needed and as needed. They are generally purpose built, as opposed to the general-purpose devices found in hyperconverged systems. A storage appliance makes storage independent from server, application, and software. This strategy allows the appliances to be deployed in amounts, places, and configurations that suit the business and the overall system, not the server or the software. It is a good solution where scalability and independence are required. A storage appliance makes it possible to add storage without adding compute and network resources. In a hyperconverged system, when expansion of storage is required, another node must be added, and a node also has compute and network resources that may not be required.

A storage appliance can be built for maximum performance, such as a database application with many users where latency must be minimized. Another application could be off site backup.

Another consideration is security. With an appliance, it is possible to absolutely determine the physical location of data. If it's stored on a NAS or SAN, it is in an identifiable piece of hardware. Data stored in a hyperconverged system is being allocated among the resources in the system. It is also easier to add security to a single device than it is to an entire hyperconverged system.

With hyperconverged systems, often the hardware has to be validated by the software manufacturer, limiting the choices of hardware. With traditional storage as an appliance, the user is free to choose the hardware that best fits the application. And since compute, network, and storage are discreet, any of the three can be taken down for maintenance, expanded, or changed without affecting the others.

In practice, IT organizations in the medical imaging field as well as other fields will probably find themselves with areas that are hyperconverged and areas that utilize traditional schemes where compute, storage, and network are not converged. To reiterate, choose the system, software, and IT strategy that suits the organization and the application. Do not get swept up in the hype and hyperbole of the "latest and greatest" and implement something that really doesn't fit.



Hyperconverged vs. Storage as an Appliance Considerations

Considerations	Hyperconverged	Storage as an Appliance
Cost of ownership	Lower for small installations	Lower for larger installations
Scaling	Scale up	Scale out
Flexibility	Less flexible	More flexible
Suitability / size of enterprise	Better for small/med. enterprise	Better for large enterprise
Customized vs. Generic	Leans towards generic	Leans toward customized
Applications	Moderate number of applications	Large number of applications
Number of users	Moderate	Large
Amount of data	Moderate	24TB or more
Performance demands	Moderately demanding	Offloads compute function to storage appliances
Changing vs. Stable environment	Volatile - dynamic environment	Stable environment
IT labor intensive	More demanding	Less demanding
IT staff	Limited - generalized	Specialized
IT stability	Subject to change	Low turnover
Security	Secure	More secure

Below is a checklist of major points to research when considering changes to your medical imaging IT strategy:

1. Is the image archive software under consideration truly, 100% vendor neutral? If not, scratch it off the list.
2. How easily (if at all) will the system under consideration integrate with my existing infrastructure? What constraints will it place on future choices of hardware and software? No one wants a “fork lift upgrade”; removing perfectly functioning equipment and replacing it is expensive and wasteful.
3. What are the storage system requirements and abilities for Backup, Failover, and Disaster Recovery? Dig deep into any vendor’s offering to make sure these concerns are well considered, and well provided for, as well as easy to deploy and manage.
4. If your current environment is virtualized, will the proposed system support existing hardware and environment? Select a vendor with devices that are supported by your existing virtualization software. If you have a VMware® environment, selecting a device that supports only Hyper-V® is not the best choice, and vice versa.
5. Maintenance and administration - look for things like a Web GUI, an intuitive interface, the ability to write scripts, and any other features that make the software easy to use.
6. Scalability: What is the limit, if any, of storage capacity? Does the system support thin and over provisioning? Does it support inline compression and data deduplication?
7. Support plans and cost. Your business will depend on this system, so excellent support at a reasonable cost is a must. Most vendors offer different levels of support – be sure there is one that fits your organizations abilities and comfort zone.

8. Any proposed system should function in all-flash, all conventional disk, or hybrid configurations. The system should support fully automated tiering. This function automatically moves the data that is most frequently accessed to SSD or flash, while storing more sedentary data on spinning disks.
9. Are there provisions for off-site backup? Snapshots? Protection against ransomware? Do not short-change yourself in these areas; spend as much money and time and effort as you think your data is worth. If you have doubts, ask someone who fell victim to the recent “wanna cry” ransomware.

About Nfina Technologies

Nfina Technologies is a US based manufacturer of Servers, and Storage products, Hyper-converged clusters and Edge Computing solutions that combines current high-performance technology with a market leading 5-year warranty & tech support. Nfina provides the best value and lowest TCO in the industry.

© 2019 Nfina Technologies, Inc. All rights reserved. Product names mentioned are for identification purposes only and may be trademarks and/or registered trademarks of their respective company.