

Converged vs. Hyperconverged IT Architectures



The Datacenter has evolved from the traditional (non-integrated) subsystems to a software-defined unified solution. Hyperconverged Infrastructure combines storage, compute, virtualization, clustering software resources, and networking resources into a single integrated system. For clustering and virtualization, Hypervisors run a multi-tasking software that distributes the application loads on Virtual Machines (VM's). Examples of vendors and implementation of this software include: VMware (Esxi), Microsoft (Hyper-V, Azure stack), Nutanix (Acropolis), and Red Hat (KVM).

HYBRID CLOUD SOLUTIONS

CYBER SECURE SERVERS

DATA STORAGE

PCs & THIN CLIENTS

BACKUP & RAPID DR

nfina
HYBRID ■ CLOUD SOLUTIONS

Converged Virtualization Architecture

The Converged virtualization infrastructure consists of clustered Hypervisors and clustered Storage and network switches as separate appliances. See Figure 1 below.

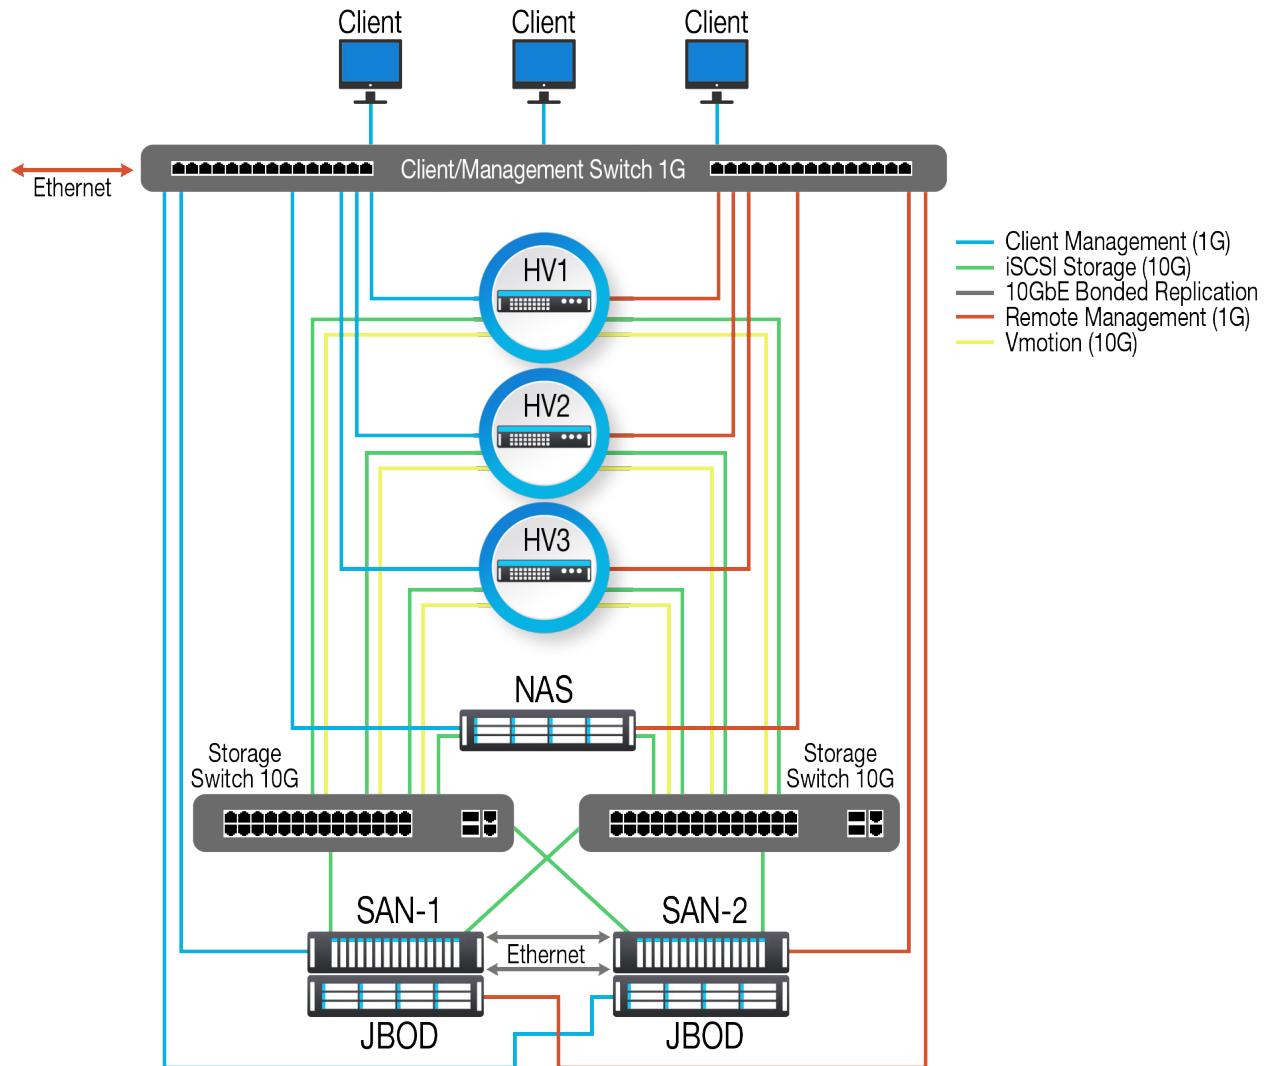


Figure 1. A typical Converged Architecture

As shown in Figure 1, there are three hypervisors (HV1-3, responsible for compute) as well as a Dual Node SAN (SAN1-2, responsible for primary storage), and a backup appliance (NAS). All of these systems have separate management interfaces and configuration requirements. In this traditional architecture, the storage and the storage network is completely separated from the Hypervisors.

This architecture is very scalable as well. If one needs more compute, you can easily add more hypervisors (up to the limit of the open ports on the switches). If more storage is needed, add more JBOD's up to 6 JBODs per SAN node as a practical limit.

For production Storage, Nfina highly recommends SAN's with 4-way mirrored SSD's for performance and high reliability. Fiber Channel is still widely used with SAN networks, but because of the expense of the network components (especially the switches), today lower cost SANs are typically configured for iSCSI block storage and SMB/CIFS file shares. NSF is still used and is fully supported by Nfina's SANs.

The NAS storage in Figure 1. holds the immutable snapshots produced by the SAN storage. Best practice for sizing: the NAS storage should be approximately twice the size of the useable production SAN array.

Hyperconverged Virtualization Architecture

Note that in the hyperconverged architecture, the storage is controlled by a virtual machine running the Nfina-Store™ storage software on the hypervisor and the storage network can be as simple as direct connect ethernet or Fibre Channel between nodes. VMware based systems running Nfina-Store utilize pass thru mode to access the storage directly from the 12G backplane, thus insuring low latency IOPs.

This solution is very cost effective, especially with the Gen3 servers now available from Intel with the higher core count CPUs. The storage is very similar to the traditional SAN arrangement (both SSD's and HDD spinners and NVMe caching are supported), as well as the iSCSI, NFS, and SMB/CIFS protocols.

This architecture is simple yet very scalable. If one needs more compute, you can easily add another hypervisor (need to also move up to a switched network using switches if a third node is added). If more storage is needed, add more JBODs. The minimum number of clustered servers required is 2 in Nfina's Hyperconverged HA solution.

For a backup pool, additional large capacity drives may be added to the hyperconverged cluster (up to the limits of the unused bays or by adding an additional hypervisor). Adding an additional hypervisor also adds more cores, which could also be an advantage when scaling. See Figure 2 below.

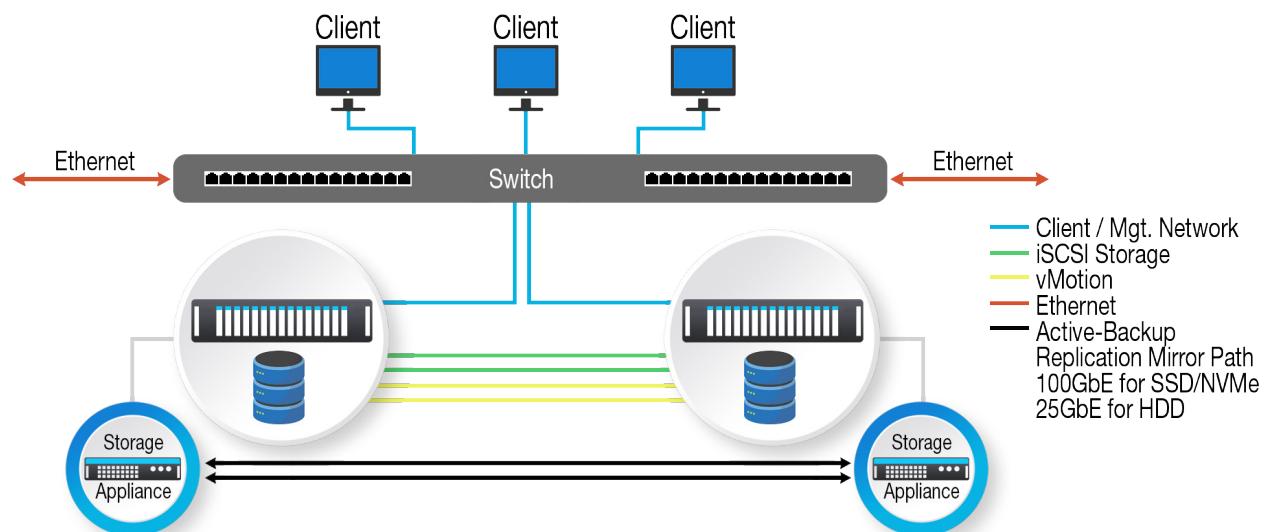


Figure 2. A typical Hyperconverged Architecture

Conclusion:

All in all, a hyperconverged infrastructure has low up-front costs, making it easier to manage capital expenses and avoid over provisioning. A hyperconverged infrastructure drastically improves scalability and agility while also enabling a software-centric architecture. To make your IT landscape even more efficient – choose Nfina's Hyperconverged. Nfina Hyperconverged offers:

- Affordable with business continuity built-in
- High-Availability Dual-Controller Clustering
- Highly configurable network options
- Intel Xeon Scalable Processors
- High-performance storage for backup, disaster recovery, or business continuity solutions

Type	Pros	Cons
Converged	<ul style="list-style-type: none"> • Highly Scalable • More Flexible 	<ul style="list-style-type: none"> • Higher Network Complexity • More Expensive • More Hardware = Higher Power Consumption • More rack-space required. Minimum implementation requires 4 nodes
Hyperconverged	<ul style="list-style-type: none"> • Scalable is good. Must add switches when more than 2 nodes are used • Simpler architecture. • Less hardware = less power consumption • Less rack-space is required. Minimum implementation requires 2 nodes. 	<ul style="list-style-type: none"> • Compute performance. Some CPU cores will be required to managed storage and networking • Special Care must be taken when performing maintenance on Hyperconverged systems due to the storage being integrated with the hypervisor.

References:

1. <https://www.vmware.com/products/hyper-converged-infrastructure.html>
2. <https://www.nutanix.com/hyperconverged-infrastructure>
3. <https://www.redhat.com/en/topics/hyperconverged-infrastructure>
4. <https://www.networkworld.com/article/3207567/what-is-hyperconvergence.html>
5. <https://www.scalecomputing.com/blog/what-is-hyperconvergence-a-complete-guide-to-hci>

© 2022 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.