

## Integrating Hyper-converged Systems with Existing SANs

By Dan Kusnetzky

### EXECUTIVE SUMMARY

Hyper-converged systems introduce an attractive new way to share internal storage among clusters of servers in a compact, cost effective configuration. The challenge these systems pose is that most deny access to valuable storage resources that the enterprise already has in place. This limitation creates data silos which can result in wasted space and a need to over provision storage overall.

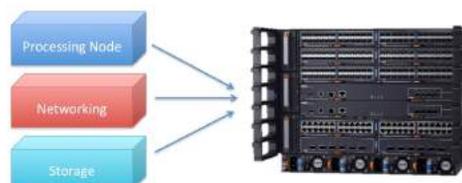
In this paper KG explores how DataCore Software integrates hyper-converged systems with existing SANs and Cloud storage to unify the data infrastructure across an enterprise, measurably reduce costs and simplify data sharing and data protection.

### THE BENEFITS OF HYPER-CONVERGED SYSTEMS

Enterprises have long segregated individual system functions like processing, networking and storage into separate devices, seeking to improve workload agility and flexibility. Unfortunately, the separation of functions increases complexity, power consumption and use of limited data center floor space.

New hyper-converged systems reduce complexity by re-integrating into a single enclosure these co-dependent functions. The compact configurations also reduce costs of administration, floor space, power consumed and heat produced.

Hyper-converged systems address growing workloads by adding more multi-function servers. This is known as a "scale-out" architecture.



### THE LIMITATIONS OF HYPER-CONVERGED SYSTEMS

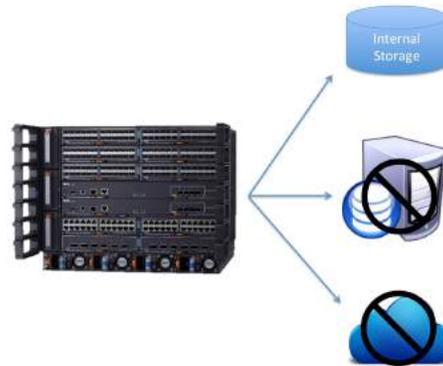
While hyper-converged systems offer enterprises many benefits, their adoption also brings along a number of limitations.

Many hyper-converged system architectures support only the use of storage that is inside of each server. There is no way to access external storage that is already in use in the data center. After enterprises have spent decades deploying shared storage environments (SANs) that support multiple hosts, this approach appears to wind the clock backward and reintroduces separate islands of storage. So, while hyper-converged systems make managing processing and networking easier, they've reintroduced a new level of storage complexity and storage management issues.

Furthermore, these systems were not designed to maintain large amounts of data locally. This means that as data accumulates some data must be moved off of these hyper-converged systems into storage systems better suited for longer retention. This can make big-data or large scale transactional workloads difficult to support.

Another challenge is that the vendors of hyper-converged systems often insist on supporting only their processing nodes, memory, networking components and storage. This limits choice, increases the costs associated with these systems and makes it impossible to tap into capacity already available in the data center.

Although the vendors typically provide proprietary tools to monitor and manage these hyper-converged systems, it may not be possible to use familiar enterprise management / monitoring tools. This can mean that IT operations staff must be re-trained or additional staff expertise must be acquired. It can also mean purchasing more storage than is really required for the workloads the enterprise has deployed.



DataCore offers hyper-converged alternatives that integrate with existing SANs and Cloud storage. In addition, it enables data to be moved from one hyper-converged system to another through the intermediary of external storage.

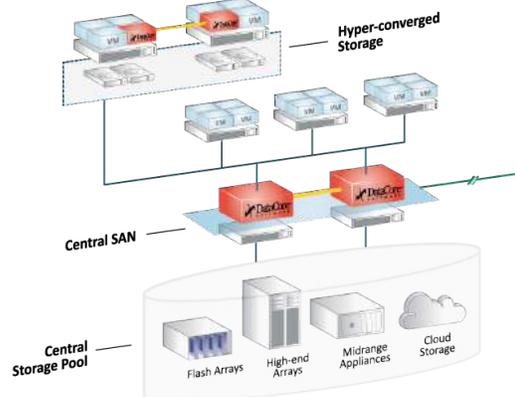
This makes it possible for the enterprise to get the best of all three approaches. In fact, the software automatically migrates data across these topologies (internal storage in the hyper-converged system, the enterprise SAN and cloud storage) to reduce costs of storage, simplify storage management; enhance data sharing between and among workloads, and make optimal use of storage space.

**INFRASTRUCTURE - WIDE APPROACH TO STORAGE**

DataCore's integrated solution provides a common set of storage services across all storage devices to make the best use of hyper-converged systems, external SANs and off-premises Cloud capacity. Storage, regardless of model or brand, can be managed as if it were a single infrastructure-wide storage resource with no wasted space.

These unified storage resources can be accessed by widely distributed applications without facing protocol or communication problems. The global shared pool of storage reduces complexity and administrative costs. Furthermore, data is replicated between storage nodes eliminating single points of failure.

In addition, workloads hosted on hyper-converged systems can easily collaborate with workloads executing elsewhere. It also makes it easier for workloads and their associated storage to be migrated into and out of hyper-converged systems.



Enterprises may scale up via larger servers and storage devices or scale out to use more servers and storage devices depending upon the performance and capacity requirements. No longer are organizations forced to select storage from a single supplier.

#### F E W E R   R E S O U R C E S   N E E D E D   –   L O W E R I N G   C O S T S   A N D C O M P L E X I T Y

Peer-reviewed, 3<sup>rd</sup> party-certified benchmarks representative of mainstream data center workloads demonstrate that DataCore markedly reduces the number of servers required to handle the most latency-sensitive applications. Much of those cost savings comes from harnessing multi-core servers to process I/O requests from multiple virtual machines (VMs) using DataCore Parallel I/O technology on a single system.

Centralized storage management spanning hyper-converged storage, external SANs and even cloud-based storage reduces the cost and complexity of allocating, protecting and monitoring storage in today's multi-faceted environment.

System administrators trained in VMware vSphere can even self-provision storage using their familiar procedures thanks to DataCore support for Virtual Volumes (VVols), even when the physical storage devices underneath don't support it.

#### S U M M A R Y   A N D   R E C O M M E N D A T I O N

While hyper-converged systems promise to reduce complexity, space and power consumption, current implementations also present challenges in the areas of storage compatibility, integration and management.

In particular, organizations need ways to complement storage inside hyper-converged systems with those housed in established storage assets, or face recreating costly islands of storage.

DataCore software-defined storage solutions address many of those challenges and shortcomings by integrating hyper-converged systems with SAN and Cloud storage. They help enterprises fully realize the combined value of their varied storage assets.

KG suggests that readers learn first-hand what DataCore can do for them. Visit <http://www.datacore.com/products/datacore-hyper-converged-virtual-san> to download a report, schedule a demonstration or download a trial version of their software.