



► Scale-out Data Management with Commvault HyperScale™ Technology

ABSTRACT

This whitepaper discusses how Commvault HyperScale™ Technology can provide the right architectural approach to build a sound data management strategy as organizations struggle with explosive data growth and adopt hybrid cloud IT.



Contents

EXECUTIVE SUMMARY 3

THE PROBLEM OF DATA GROWTH AND LEGACY DATA MANAGEMENT ARCHITECTURES 3

INTRODUCING COMMVAULT HYPERSCALE™ TECHNOLOGY 4

COMMVAULT HYPERSCALE™ TECHNOLOGY ARCHITECTURE OVERVIEW 4

Scalability and Resiliency 5

COMMVAULT HYPERSCALE™ TECHNOLOGY DEPLOYMENT CHOICES 6

Commvault HyperScale™ Software 6

The Commvault HyperScale™ Appliance 6

KEY BENEFITS OF COMMVAULT HYPERSCALE™ TECHNOLOGY 7

UNIQUELY BROAD PLATFORM SUPPORT 8

CONCLUSION 9



EXECUTIVE SUMMARY

As organizations continue to innovate their businesses through digital transformation, it is clear that data has become the new currency. Successful organizations must harness the power of data to drive competitive differentiation and provide value to their customers.

Managing the lifecycle of data across primary, secondary, and tertiary tiers in today's hybrid IT environments represents a daunting challenge for organizations. Based on a study conducted by Commvault, more than 70% of data now exists in the secondary storage tier¹. With such data growth in the secondary storage tier, it is clear that previously used approaches for data management are no longer sufficient. It is also clear that while some workloads may be suitable for cloud migration, substantial workloads will remain on premises. Therefore, providing cloud-like scale, resiliency, and economics is top of mind for many organizations.

Commvault HyperScale™ Technology is a modern approach to data management that leverages the breadth and depth of Commvault's data management platform combined with the benefits of a scale-out architecture. Available in a software or a turnkey appliance form factor, Commvault HyperScale™ Technology offers a modular, scalable, and highly resilient data management solution that is simple to install and manage. Commodity hardware and lower operational demands, combined with subscription-based pricing; allow customers to enjoy cloud-like economics while reducing the overall TCO of their data management solution.

THE PROBLEM OF DATA GROWTH AND LEGACY DATA MANAGEMENT ARCHITECTURES

Over the past few years, a lot has been written about the explosive data growth organizations are experiencing as they go through digital transformation.

Data management solutions over the past two decades have been based on a traditional scale-up architecture. As data has grown, a continuous cycle of fine tuning this monolithic scale-up architecture is a requirement – demanding time and resource complexity to manage. However, the basic architecture has largely remained the same and has consisted of media servers playing the role of data movers with dedicated storage capacity in the form of HDDs at the back end or storage sub-systems, which provide higher storage capacity and performance with the addition of storage expansion. As these media servers still largely govern performance in terms of IOPS and throughput by adding storage or scaling-up, performance saturation is reached at the compute layer once capacity starts increasing. This performance slow down often means missed SLAs, and the business being exposed to risk. Additionally, these legacy


Commvault Hyperscale Technology Solution Brief

Enterprise IT is being transformed with the maturing of public cloud providers that offer compute, storage and application services with unprecedented elasticity, scale, resiliency and availability, on a consumption based economic model.

READ NOW



commvau.lt/2Fm9FeE



architectures were built on disparate pieces of infrastructure which are generally managed by different teams within the organizations, which created several dependencies and delays in getting a solution into production. Management, patching and upgrades were considered to be equally onerous tasks. In multi-tenancy and shared service environments this task complexity is multiplied many times over based on the need for change control, end customer communication, and highly constrained operational windows.

► INTRODUCING COMMVAULT HYPERSCALE™ TECHNOLOGY

The combination of hyperconvergence and distributed web-scale computing architectures is revolutionizing enterprise datacenters. With hyperconvergence, organizations are able to collapse IT infrastructure silos and bring cloud-like agility and economics to on-premises IT operations. Web-scale extends the concept of hyperconvergence using a highly distributed, shared-nothing architecture, with features such as data availability, linear scalability of compute and storage, and non-disruptive in-place upgrades.

Commvault HyperScale™ Technology is a hyperconverged data management solution built upon Commvault's industry leading technology, Commvault Complete™ Backup & Recovery, that tightly integrates compute, storage, full life cycle data management, and analytics into a single platform across the data center and the cloud. It includes storage and server compute (CPU and memory) in a single platform building block. Each building block consists of three nodes based on industry-standard and high performance x86 server technology and delivers a unified, scale-out, shared-nothing architecture with no single point of failure.

Commvault HyperScale™ architectures consolidate all the roles performed by discrete servers in a traditional data management architecture into a single software defined stack. The software spans multiple nodes running on general purpose x86 servers, creating a massively addressable storage pool with built-in enterprise class data management capabilities. Commvault HyperScale™ Technology eliminates the need for dedicated media servers, proprietary controller based storage devices and cloud gateways, dramatically reducing complexity and infrastructure costs.

Commvault's bespoke cloud-ready data management capabilities allow Commvault HyperScale™ Technology users to transform a significant portion of their enterprise secondary storage infrastructure to provide cloud-like simplicity, elasticity, resiliency, flexibility and scale for secondary use cases. Commvault HyperScale™ can meet the most demanding RPO/RTO and instant data access needs by utilizing Active Copy Management, hardware snapshot integration and support for all major enterprise applications, databases, and cloud platforms. A sophisticated orchestration engine, programmable API, and direct data access through standard interfaces allows users to embed these data management services, seamlessly into a broader cloud framework. With low hardware acquisition costs and simplified operations, organizations can build a true cloud-based data management service for the modern hybrid datacenters.

► COMMVAULT HYPERSCALE™ TECHNOLOGY ARCHITECTURE OVERVIEW

Commvault HyperScale™ Technology is a highly scalable and resilient data management solution. A typical HyperScale deployment starts with a 3 node block. The nodes consists of standard x86 servers with built-in storage that consists of Hard Disk Drives (HDDs) for hosting backup data and Solid-State Drivers (SSDs) for hosting operations that require high performance storage. The HDDs from each server are aggregated into a scale-out clustered storage pool that is available as a single mount point from any of the nodes in the deployment. There are two kinds of SSDs recommended within the nodes. SATA SSDs are used for hosting the operating system and the

Commvault binaries on each of the nodes. NVMe-based SSDs are used for the deduplication database and the index cache. Commvault HyperScale™ supports inline source side deduplication which reduces network load during data movement operations. Commvault deduplication is content aware, policy-based, and tunable to specific workloads to minimize impact to production workloads. The deduplication database is spread across multiple partitions allowing for continued operations even with loss of multiple nodes.

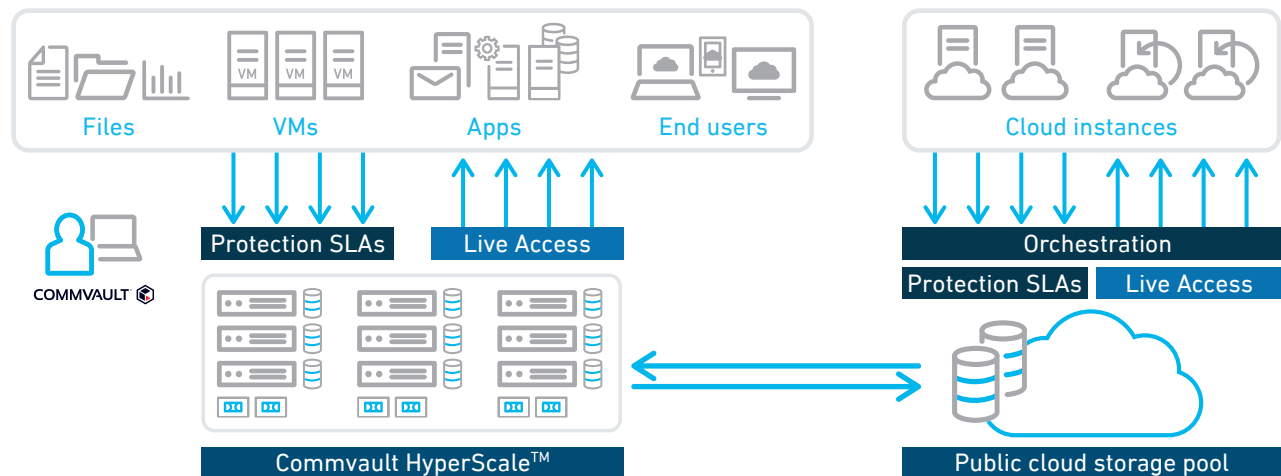


Fig 1: High Level Architecture of Commvault HyperScale™

SCALABILITY AND RESILIENCY

Deployment of Commvault HyperScale™ Technology can be easily scaled with additional blocks. Such expansion adds predictable compute and storage capacity and can be done in-line with no downtime. The data on the existing nodes is automatically redistributed on the new nodes to maintain optimal capacity and performance on the entire cluster. Users can also to mix and match multiple generations of hardware for maximum flexibility.

Commvault HyperScale™ Technology also provides enterprise grade resiliency with protection built in at various levels. Within the Commvault HyperScale™ Technology architecture, each node plays the role of a MediaAgent, which is responsible for data movement and processing of the backup data. If a node fails, jobs are automatically routed to other nodes in the cluster, as defined in the backup policy. The data itself is protected via erasure coding. Erasure coding (EC) is a method of data protection in which data is broken into fragments, expanded and encoded with redundant data pieces and stored across a set of different locations. This allows the recovery of the data stored on one or more nodes in case of failure. Various erasure coding combinations are supported which allow customization of data resiliency based on the organization's policy.

COMMVAULT HYPERSCALE™ TECHNOLOGY DEPLOYMENT CHOICES

Commvault HyperScale™ offers deployment flexibility with the following choices:

- 1 **Commvault HyperScale™ Software:** software solution that can be implemented on a hardware platform of the user's choice. Also available as validated reference designs with leading hardware providers. Ideal for larger deployments that are looking for custom configurations or to optimize for density, performance, and economics.
- 2 **Commvault HyperScale™ Appliance:** turnkey appliance based solution built, sold, and supported directly by Commvault. Ideal for medium sized environments that are looking to eliminate the time and effort of installing, configuring, and maintaining their data management infrastructure. Also an integral part of larger enterprises for ROBO-based protection / data management workloads.

COMMVAULT HYPERSCALE™ SOFTWARE

Commvault HyperScale™ Software consolidates all the roles performed by discrete servers in the traditional data protection architecture into a single software defined stack. The software spans multiple storage nodes running on general purpose servers, creating a massively addressable storage pool with built-in enterprise class data management capabilities. This eliminates the need for dedicated media servers, proprietary controller-based storage devices and cloud gateways, reducing infrastructure costs dramatically.

Customers can deploy Commvault HyperScale™ Software on a hardware platform of their choice or choose to implement it using one of the Commvault validated reference designs built in collaboration with leading hardware vendors.

COMMVAULT HYPERSCALE™ APPLIANCE

The Commvault HyperScale™ Appliance is a complete turnkey solution built and supported by Commvault. The solution tightly integrates compute, storage, scale-out file system, and the Commvault Complete™ Backup and Recovery to offer full lifecycle data management and analytics in a single platform across data centers and the cloud.

The appliance form factor saves time in hardware acquisition, installation and integration, daily management, as well as patching and updating. The single call support from Commvault simplifies the technical support process on the entire solution. The Commvault HyperScale™ Appliance offers similar scaling and resiliency benefits as described above. The fundamental architecture remains the same as a software only deployment of HyperScale.™ One added benefit of the appliance form factor is that CommServe® can be hosted natively within the appliance on the shared hypervisor layer. CommServe® runs within a Windows VM on a node of the user's choice. In case of a failure, CommServe® VM is failed over to another node within the HyperScale™ deployment, which ensures the backups continue to run.

Commvault HyperScale Appliance

Deliver seamless scale and resilient data protection with the Commvault HyperScale™ Appliance – an integrated Commvault backup appliance built on Commvault HyperScale™ Software and a scale-out infrastructure.

VISIT NOW



commvault.com/appliance



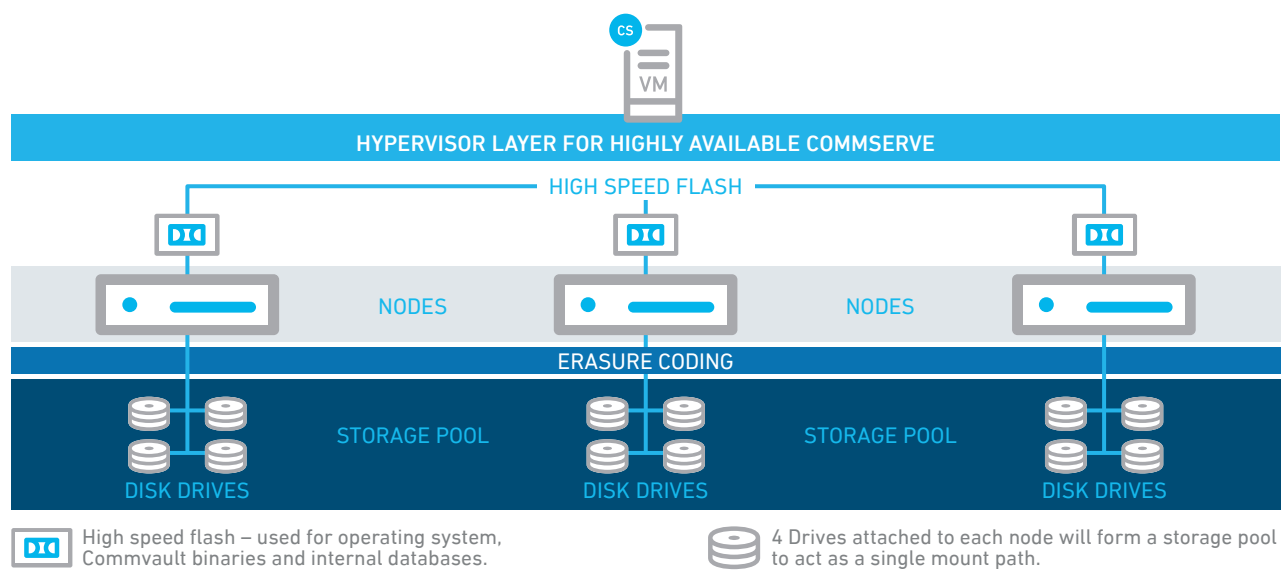


Fig 2: The Commvault HyperScale™ Appliance High Level Architecture

The appliance is available in capacity points ranging from 48TB to 432TB based on raw capacities per 3-node block. Usable capacity based on erasure data coding is approximately 66% of raw capacity.

► KEY BENEFITS OF COMMVULT HYPERSCALE™ TECHNOLOGY

Simplicity: Deploy and expand a hyperconverged pool in as little as 30 minutes. Single, easy to use management console, Commvault Command Center™, delivers simplified management of the entire solution.

Lower Infrastructure Costs: Runs on general purpose x86 servers with built-in storage. Eliminates the need for expensive external storage. Eliminates dedicated media servers further reducing hardware costs.

Resiliency and Availability: Highly resilient architecture with redundant hardware components. Erasure encoding ensures data is available through multiple drive or node failures. Partitioned deduplication and Gridstor® enable data management operations to run uninterrupted. WAN optimized replication and geo-dispersed clustering options are available for DR.

Active Copy Management: Active Copy Management for all critical enterprise applications on a variety of physical and virtual platforms for very low RPO/RTO. Tight integration with all major primary storage vendors for snapshot based copy management including new flash based devices.

High Performance Deduplication: Large deduplication pools with distributed processing across multiple nodes to meet the most demanding RPO/RTO needs.

Operational Efficiency: No fork lift upgrades needed. Evergreen storage pool with nodes that can be upgraded/replaced/fixed without disruption of services.

Scale and Flexibility: Storage pool can start small and expand dynamically as needed, up to petabytes of usable capacity. Mix multiple generations of hardware in a single pool to rapidly benefit from newer architectures and drive densities.

Instant Data Access: Restore-less access to data copies by users and applications using standard interfaces from all managed copies including cloud copies. Scale-out compute nodes can handle the most demanding read requests across several users or applications.

Cloud Integration: Tier to all major public cloud storage providers for offsite copies and long term retention. Transform on-premises workloads into public cloud instances, protect cloud-native workloads and replicate data back to on-premises. Provisioning policies to spin up/spin down cloud resources on demand.

Automation and APIs: Programmable workflow engine to automate the most complex tasks involving the usage of secondary storage copies, and create custom portals to suit the unique needs of an organization. Command line tools and REST API available to be used with multiple programming tools including C#, Python and Ruby.

Validated Reference Designs: Validated reference designs with leading hardware providers to offer flexibility and choice for underlying hardware.

Simplified Support: Commvault® is responsible for supporting all elements of the software stack including the operating system for the storage pool nodes. On the the Commvault HyperScale™ Appliance, complete stack support, inclusive of hardware, is handled directly through Commvault.

Flexible Pricing Options: Simple, flexible and subscription based pricing models for cloud like economics and quicker ROI.

The Commvault HyperScale™ Appliance Solution Brief

The Commvault HyperScale™ Appliance is a hyperconverged data management solution built upon Commvault's industry leading technology that tightly integrates compute, storage, full life cycle data management and analytics into a single platform across the data center and the cloud.

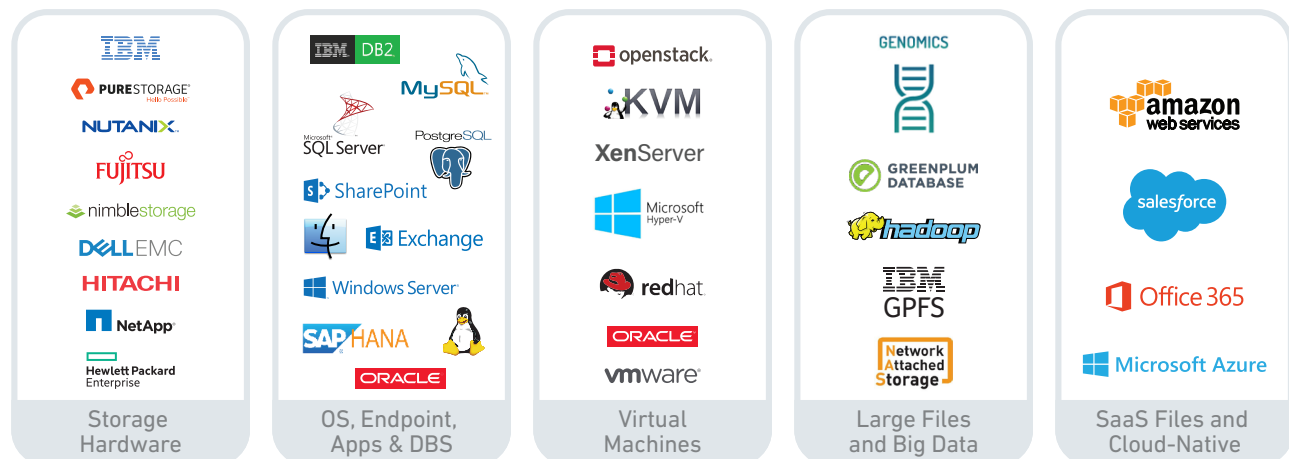
READ NOW



commvau.lt/2DKkW71

► UNIQUELY BROAD PLATFORM SUPPORT

Commvault – a Gartner Magic Quadrant Leader in Data Protection² – understands that when it comes to data management on-premises and in the cloud, customers have to deal with various kinds of workloads. With that in mind, Commvault HyperScale™ Technology supports over 200 platforms which includes various operating systems, applications & databases, storage platforms, big data, SaaS and cloud native applications to provide comprehensive coverage for workloads in the industry.



COMMVault INTEGRATION



► CONCLUSION

Legacy methods of data management are no longer sufficient to meet the data management needs arising due to explosive data growth. Commvault HyperScale™ Technology provides a modern data management solution that is scalable, highly resilient, and simple. It allows organizations to manage data seamlessly on-premises and in the cloud. Commvault HyperScale™ Technology removes the burden of day-to-day operations with a simplified installation, automated and self-service operations, and update process. Available as a software solution or a turnkey appliance, based on a subscription licensing model, Commvault HyperScale™ Technology allows customers to modernize their data management operations with cloud-like availability, scale, and economics.

► Evolve your on-premises environment so that it mirrors the business benefits you've achieved using cloud. Learn more at commvault.com/hyperscale.

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

© 2018 Commvault Systems, Inc. All rights reserved. Commvault, Commvault and logo, the "C hexagon" logo, Commvault Systems, Commvault OnePass, CommServe, CommCell, IntelliSnap, Commvault Edge, and Edge Drive, are trademarks or registered trademarks of Commvault Systems, Inc. All other third party brands, products, service names, trademarks, or registered service marks are the property of and used to identify the products or services of their respective owners. All specifications are subject to change without notice.

