

Commvault HyperScale

Debrief

Table of Contents

Secondary Storage EcoScape.....	1
Commvault Hyperscale: Introduction.....	2
Commvault Hyperscale: For the Technical Staff	3
Fujitsu as a Server Partner.....	4
Software Architecture.....	4
Commvault's Software	5
Cloud Integration.....	5
The Business Side of Commvault Hyperscale	6
Summary.....	7

Secondary Storage Landscape

The secondary storage market has emerged in recent years as an increasingly interesting conglomeration of the storage and data protection markets. While this combination may not have been intended to launch the modern iteration of secondary storage, it's clear that data protection takes a front-and-center seat as the primary use case in secondary storage. Although certain solutions go beyond the data protection use case, for many organizations, data protection defines the full breadth of needs around non-primary data storage.

Most organizations have implemented robust primary storage systems that feature the characteristics necessary to operate their key business workloads. Oftentimes, these primary storage systems require high levels of I/O for certain workloads. Such systems typically have strict service level agreements attached to them as well. Because of the robustness of the performance characteristics and the need for a business-critical set of SLAs, these types of solutions sometimes carrying significant price tags that can be off-putting for some, particularly when there is a mix of workloads, with some not requiring the level of robustness as a mission critical one.

But, primary storage accounts for just a fraction of an organization's total storage footprint. As mentioned, not every workload needs 99.999% uptime and minimum IOPS guarantees. There are a vast number of workloads that simply don't need primary storage performance characteristics. There is also an increasing abundance of

Who is Commvault?

Commvault was incorporated in 1996 after being spun out of Bell Labs and AT&T. Listed on the NASDAQ (CVLT), Commvault is headquartered in Tinton Hills, NJ and, in 2016, had revenue approaching \$600 million.

Commvault is recognized by Gartner as a leader in its Magic Quadrant and has been a consistent leading provider of data protection solutions.

About the Author

Scott D. Lowe is CEO of ActualTech Media, where he leads content and product strategy. Prior to co-founding ActualTech Media, Scott spent ten years as a CIO for various organizations. He rose through the ranks, spending time as a developer, a network engineer, a systems engineer and more along the way. With a wide breadth of experience, Scott brings insight to both technical and decision-making audiences. In addition, Scott is still a practicing consultant, providing strategic IT and business direction to his clients.

© Copyright 2017, ActualTech Media

non-primary data that needs to be retained for business, legal, or regulatory reasons. Purchasing tier 1 level storage for all storage needs, including data protection, can be prohibitively expensive—and unnecessary. The acceleration in data growth as created the need for full-featured data reduction technologies—deduplication and compression—to be included in storage systems of all shapes and sizes. To this point, many emerging storage systems, both primary and secondary, feature data reduction

At the same time as this tug-of-war is taking place between primary and secondary storage, organizations are looking for ways that they can reshape their storage environments to more closely resemble public cloud environments. This desire for rapidly deployable scale-out storage solutions is driving interest in new kinds of storage technologies and has given rise to a number of new architectural opportunities, including hyperconverged infrastructure and other server-based solutions.

The conglomeration of all of these challenges and trends has led to the development of a number of approaches to attacking the secondary storage challenge.

Commvault Hyperscale: Introduction

Recently, Commvault jumped into the secondary storage market with the introduction of their scale-out Hyperscale product. Hyperscale is deployed on a series of x86 servers, although there are opportunities to buy from other partners—such as Cisco, sold under the product name ScaleProtect, which we'll discuss shortly—as well. There are a number of reasons that Commvault has ventured into this space, not the least of which is the fact that there are compelling startups that have been quite successful with their own secondary storage offerings. Cohesity and Rubrik are two that are immediately obvious and both feature data protection as a primary use case, although Cohesity goes well beyond to embrace other needs as well. In addition, even established storage players, such as HPE's Nimble unit, have introduced secondary storage offerings. In the case of Nimble, the company partners with Veeam to provide a complete data protection hardware and software combination.

Commvault HyperScale: For the Technical Staff

Commvault's HyperScale product can be acquired in one of two ways. First, customers can buy a multi-node solution directly from Commvault. Second, if the Commvault-provided product is not sufficient to meet data protection needs, Commvault has developed a reference architecture option that any server vendor can adopt and build a comprehensive and fully supported data protection solution. As of this writing, Commvault has developed a deep relationship with Cisco and also has partnerships with Lenovo, HPE, Super Micro Computer, Huawei and Dell-EMC.

The Commvault-provided solution offers a true pay-as-you-go appliance-based data protection service. By deploying a minimum of three HyperScale nodes, you're essentially deploying a complete secondary storage environment that can withstand the loss of a node or multiple drives in nodes. This three-node configuration is referred to as a *block* by Commvault and every node in the block is identical. You can add additional blocks, but within each block, the hardware needs to remain consistent.



Figure 1: Commvault HyperScale blocks consist of three individual nodes

Each HyperScale node features raw capacity of 16, 24, 32, or 40 TB. With a three-node minimum configuration, customers get a block that includes 48 TB to 120 TB of raw capacity. Further, each node includes 150 GB of M.2-based flash storage, or 450 GB for each block. This hybrid configuration enables the HyperScale solution to provide both performance and capacity. Each node features 96 GB of RAM, an 8-core Xeon processor from Intel, and dual 10 Gb Ethernet adapters.

The capacities listed above are raw. When combined with data reduction services, such as deduplication and compression, it's likely that customers will see much higher effective capacity.

Fujitsu as a Server Partner

The Commvault HyperScale solution is made available via Commvault's new partnership with Fujitsu, the server vendor on which the solution operates. I spoke with Commvault about their decision to work with Fujitsu, and their reasoning is quite good. First, Fujitsu was very eager to work with Commvault, who preferred to work with a server vendor that would not impose a massive minimum purchase. With HyperScale being a new offering from Commvault and something of a departure from the company's software-only roots, there was some question about what sales might look like. In the unlikely event that sales don't go where expected, Commvault didn't want to be on the hook for a huge minimum order.

But that's actually not as critical as ongoing logistics. Commvault could have engaged a company like Supermicro, but there were concerns about ongoing support. Commvault needed a partner that eliminated worries about supply chain and parts depots. Remember, Commvault is providing direct worldwide support on both the hardware and the software, and they needed a hardware partner that could guarantee such support.

By partnering with Fujitsu, Commvault gets instant access to Fujitsu's supply chain and parts depots. The beauty of the arrangement is that Fujitsu will provide technicians. Commvault just has to do some basic log analysis and send those logs over to Fujitsu when a hardware problem is identified. For Commvault, this resulted in a very low barrier to entry and minimal risk.

Trying to build out a supply chain from nothing is an expensive and long-term effort. It can take years to recoup the initial investment. Via the Fujitsu relationship, Commvault has not had to invest anything but time and integration and they have been able to go to market with a fully supported solution far more easily and cost-effectively than would otherwise be possible.

Software Architecture

HyperScale is, for all intents and purposes, a storage platform designed specifically to support Commvault. Today, organizations need the ability to very easily scale both their primary and secondary storage environments. To provide the level of scalability needed for HyperScale, Commvault partners with Red Hat and implemented HyperScale on Red Hat's Gluster Storage. Gluster provides the platform with comprehensive scale out functionality, load balancing, and data reduction. Availability is achieved through Gluster's erasure coding services, and allow the platform to lose multiple disks or even a full node and remain available.

I asked Commvault why they chose Red Hat. In short, it comes down to viability and economics. The company has no interest in being a storage company. They want to get to the good stuff, which is managing and protecting data, but they still wanted to offer an appliance-based solution. Gluster is already built and already proven. Commvault chose the Red Hat distribution because Red Hat is the biggest contributor into the Gluster community. As a result, Commvault gets more ability to inform direction to Gluster.

Commvault's Software

On top of this hardware and scale-out storage software pyramid sits Commvault. No more do customers need to go through the effort of deploying Commvault's software. It comes pre-installed and is ready to go. In addition, Commvault provides updates to the platform on a regular basis, at least quarterly. The appliances can be upgrade in place.

Cloud Integration

Data protection has become a key use case for companies consider cloud initiatives. The lack of a desire and funds to build a secondary data center has pushed many an organization into the arms of Amazon, Microsoft, and others to meet this secondary need. In many ways, cloud is replacing tape in terms of meeting long-term data archival needs. Cloud is eminently scalable and, when used appropriately, cost-effective. More importantly, though, recovery can be far faster from a cloud environment, depending on the scope of the recovery, particularly when compared to recovering from tape. In addition, data can be moved into/out of the cloud or between clouds, an important step forward as companies adopt multi-cloud strategies.

Moreover, as companies like Commvault embark on the wider-ranging data management mission, as opposed to simple data protection, having the ability to manage the full lifecycle from primary to backup to disaster recovery to archive, whether in tape or cloud, is criticalThe Business Side of Commvault Hyperscale

The Business Side of Commvault Hyperscale

We're now living in the "as-a-service" world in which companies of all shapes and sizes are adopting pay-as-you-go services that impact the operational budget rather than continually requiring you to spend large capital budgets. We can thank cloud service providers for bringing this budgeting paradigm to the forefront. Over time, as so many other companies have done or are doing, you should expect to see Commvault pivoting more toward subscription-based pricing, with Hyperscale as an initial foray into that space. As the cloud has proven, organizations want to move to OpEx-driven purchases rather than continue with CapEx-intensive outright purchases. Today, customers are able to buy Commvault software as a perpetual license or on a subscription model. There are currently both models in the company's price book. That shift will not happen overnight. To help ease the transition, there is a program for people with perpetual licenses that want to switch to a subscription model.

Commvault's Hyperscale product pricing, as you might expect from an as-a-service-type offering, includes all software upgrades as well.

If you're considering Hyperscale and you're not sure if you want to buy directly from Commvault or purchase via a partner such as Cisco, there are some things you need to know. At first, it may seem like you can buy the same solution either directly from Commvault or from Cisco, but there are some really important differences. First, only Commvault is making available the true pay-as-you-go offering. When you buy via Cisco, Commvault is not making that option available. That said, you may have some creative financing options available to you via Cisco, but, as of this writing, the Commvault-direct payment model is not available.

As mentioned, only the Commvault-provided appliance enables a subscription purchasing option. As a part of the subscription, every three years, Commvault replaces all of the hardware as a part of the ongoing refresh.

Commvault's own offering is technically expandable to whatever levels you might need, but configurations top out at 100 TB of capacity. This isn't a hard limit; if you want, you can expand further, but Commvault encourages customers needing more than 100 TB to consider working with a partner, such as Cisco.

At first, this seems to be a curious limitation. After all, Commvault's offering offers a compelling purchasing option.

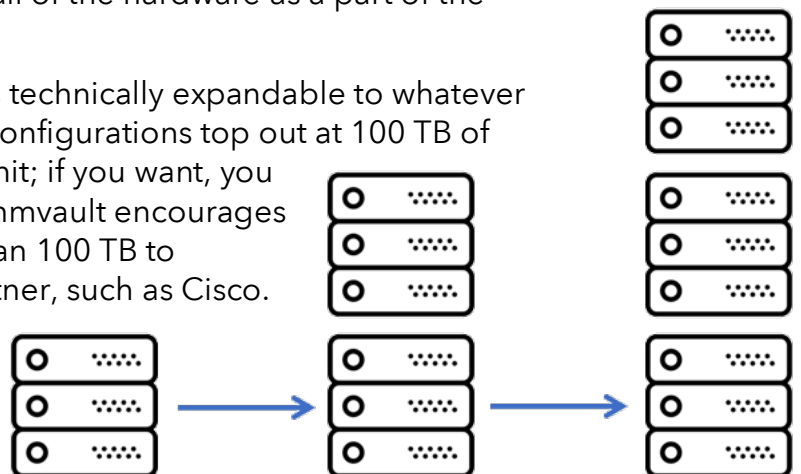


Figure 2: HyperScale nodes are linearly scalable

The 100 TB suggested limit was very intentional. It creates a boundary after which customers need to work with a partner. Although Cisco is the first partner in the HyperScale product line, it's expected that a number of other players will emerge. Under Cisco's arrangement with Commvault, the Cisco sales force is brought to bear on deals. Cisco is selling Commvault bundled with their UCS servers and Cisco sales reps get paid on the entire deal. This is not a meet in the channel approach but is a complete solution. As Commvault adds more partners to the program, they will enjoy sales force multiplier effects. To differentiate the two offerings, Cisco's version of the solution is known as ScaleProtect.

In reality, the Commvault-provided appliance exists primarily to close a gap. Competitors are selling data protection services in appliance form factors. Customers have proven that they are willing to pay for the convenience of an appliance. In fact, Commvault's research finds that most midmarket customers prefer appliances. They want to get out of the business of integrating solutions together. Commvault defines the midmarket as companies with 200 TB or less of backup storage needs. Enterprise customers are defined as those requiring above 200 TB of capacity.

Without their own appliance, Commvault risked not having a seat at the customer table. Commvault feels that the combination of their appliance and partner program provides a comprehensive offering that meets the needs of a wide variety of potential customers. With the deployment question out of the way, Commvault can talk about what they can do with the solution. People are looking for support for physical, virtual, and cloud platforms. With support for over 200 platforms and applications, Commvault believes they have a powerful and compelling message for potential customers.

It's likely that partners will eventually adopt their own pay as you go models if they don't already have something available.

Summary

As an opening product in the secondary storage space, HyperScale brings Commvault current with a number of other competitors in the market while also helping the company capitalize on the cloud movement. With both a subscription offering directly from Commvault and partner programs for particularly large implementations, Commvault aims to hit all market segments at once. My main concern with the program is the gap between where ScaleProtect is supposed to top out (100 TB) and what Commvault considers an enterprise-sized environment (200 TB). It's technically possible to go beyond 100 TB, but it's not the solution's target market.