

REFERENCE ARCHITECTURE

# FlashBlade with Commvault Cloud Design Guide

Design considerations for deploying Commvault and Pure Storage FlashBlade on Cisco UCS C220 M6.

# Contents

<b>Introduction</b>	3
<b>About This Guide</b>	3
<b>Reference Architecture</b>	4
Overview	4
Core Solution Components	7
Customizable Solution Components	7
<b>Designing for Your Environment</b>	8
Single-chassis Base Configurations	8
Customizable Components	10
Cisco UCS C220 M6 Part Numbers	10
Policy Considerations	12
<b>System Placement and Site Design</b>	12
General Guidelines	12
Multiple Sites	12
<b>Expansion Guidelines</b>	14
Predicting Future Needs	14
What to Add	14
<b>Conclusion</b>	15
Additional Resources	15
About the Author	15



## Introduction

As IT budgets continue to lag behind data growth, it has become necessary for organizations to find efficiency wherever they can. Having to invest significant time designing, testing, purchasing, implementing, and supporting a solution steals time and focus from your strategic goals. A validated solution lets you put more effort toward priorities that bring more value.

It's with this in mind that Pure Storage® and Commvault created this joint reference architecture. It provides a blueprint for deploying Commvault software on Cisco UCS hardware and storing backup data on Pure Storage FlashBlade®. This solution provides several key benefits: It simplifies the purchasing and deployment processes, saving valuable time. It makes scaling simple as your environment grows. The Pure Storage Evergreen® model means never having to migrate to a new backup storage target. Its physical density reduces total cost of ownership, with complete management over a petabyte of backup data in nine rack units. Validated designs give you confidence of a successful outcome.

Commvault and Pure Storage enhance your cyber resiliency position as well. Hardened Linux servers reduce the attack surface of your backup environment. Commvault has numerous ransomware detection and mitigation features available, including anomaly detection and Metallic ThreatWise, that reduce the chance of an attacker going undetected long enough to compromise your key data. Pure Storage FlashBlade and object lock keep an attacker from destroying your backup data.

In the event an attacker does compromise your systems, the solution will help get you back online in less than the days or weeks you've seen in ransomware headlines. The architecture can scale to support recovery of tens to hundreds of terabytes per hour. You can leverage the optional Commvault Threat Scan and Clean Room Recovery to prevent reinfection during the process.

---

## About This Guide

This guide is intended as a reference to support IT administrators designing a Commvault Cloud software deployment using the Cisco Unified Computing System (UCS) C220 M6 server and Pure Storage FlashBlade. You can use this reference architecture in a new Commvault environment or to modernize an existing one.

**NOTE:** Some links to Commvault documentation require a customer account to access them.



## Reference Architecture

### Overview

IT organizations face constant challenges. You grapple with increased risk from ongoing and emerging cyber threats. You have to handle the day-to-day operations of your environment while managing the lifecycle of your systems. You have to buy and deploy new solutions, and that often means figuring out for yourself what hardware you need.

This reference architecture addresses all of these challenges head-on with three components: the Commvault Cloud platform, backup storage, and servers to manage data movement. A prescribed bill of materials (BOM) simplifies acquisition.

A standardized configuration and automated configuration elements make deployment easier and gives you a known outcome, with enough flexibility to fit your unique needs. Commvault policy-driven management and simplicity from Pure Storage reduce the management burden. Non-disruptive upgrades and Evergreen Storage keep hardware refreshes from distracting you from strategic goals. Secure configuration and ransomware mitigation reduce your risk.

For cost-effective long-term retention, you can add FlashBlade//E™ for efficient on-premises storage with fast recovery, and optional Commvault [Air Gap Protect](#) gives you an isolated, cloud-based data copy.

The solution leverages FlashBlade fast object storage to simplify scaling and performance. A single object bucket can service all traditional retention requirements. You can configure additional buckets with object lock enabled to provide immutability against malicious attackers or to serve data isolation and other advanced retention needs.

The solution is designed to take full advantage of [Commvault Storage Accelerator](#). This optional feature lets Commvault backup agents send data directly to and from FlashBlade and bypass the MediaAgents. Direct communication makes better use of your network bandwidth and lets the MediaAgents manage more backup jobs. However, the solution delivers high performance and simplicity even without Accelerator.

To simplify deployment and administration, the architecture takes a modular building-block approach. You can start with specific combinations of C220 and FlashBlade//S200 based on your current size. As your data grows, you can expand the FlashBlade transparently, and you can add C220 servers as needed to manage the extra capacity.

Adding FlashBlade//E for long-term storage or Air Gap Protect for cloud-based cyber resilience is simple. Both are easy to configure and add, and neither requires any changes to the MediaAgents.

Figure 1 shows the components and logical architecture of the solution. The Commvault MediaAgents, as well as backup agents enabled for Storage Accelerator, use the S3 protocol to store backups to FlashBlade//S™ for rapid recovery. The MediaAgents use Commvault's efficient Auxiliary Copy to replicate new, unique data to the optional long-term recovery and cloud tiers on FlashBlade//E and Commvault Air Gap Protect or other platforms. The short- and long-term tiers can reside in the same site, separate sites, or any mix, and you can have more than one target in any tier.



REFERENCE ARCHITECTURE

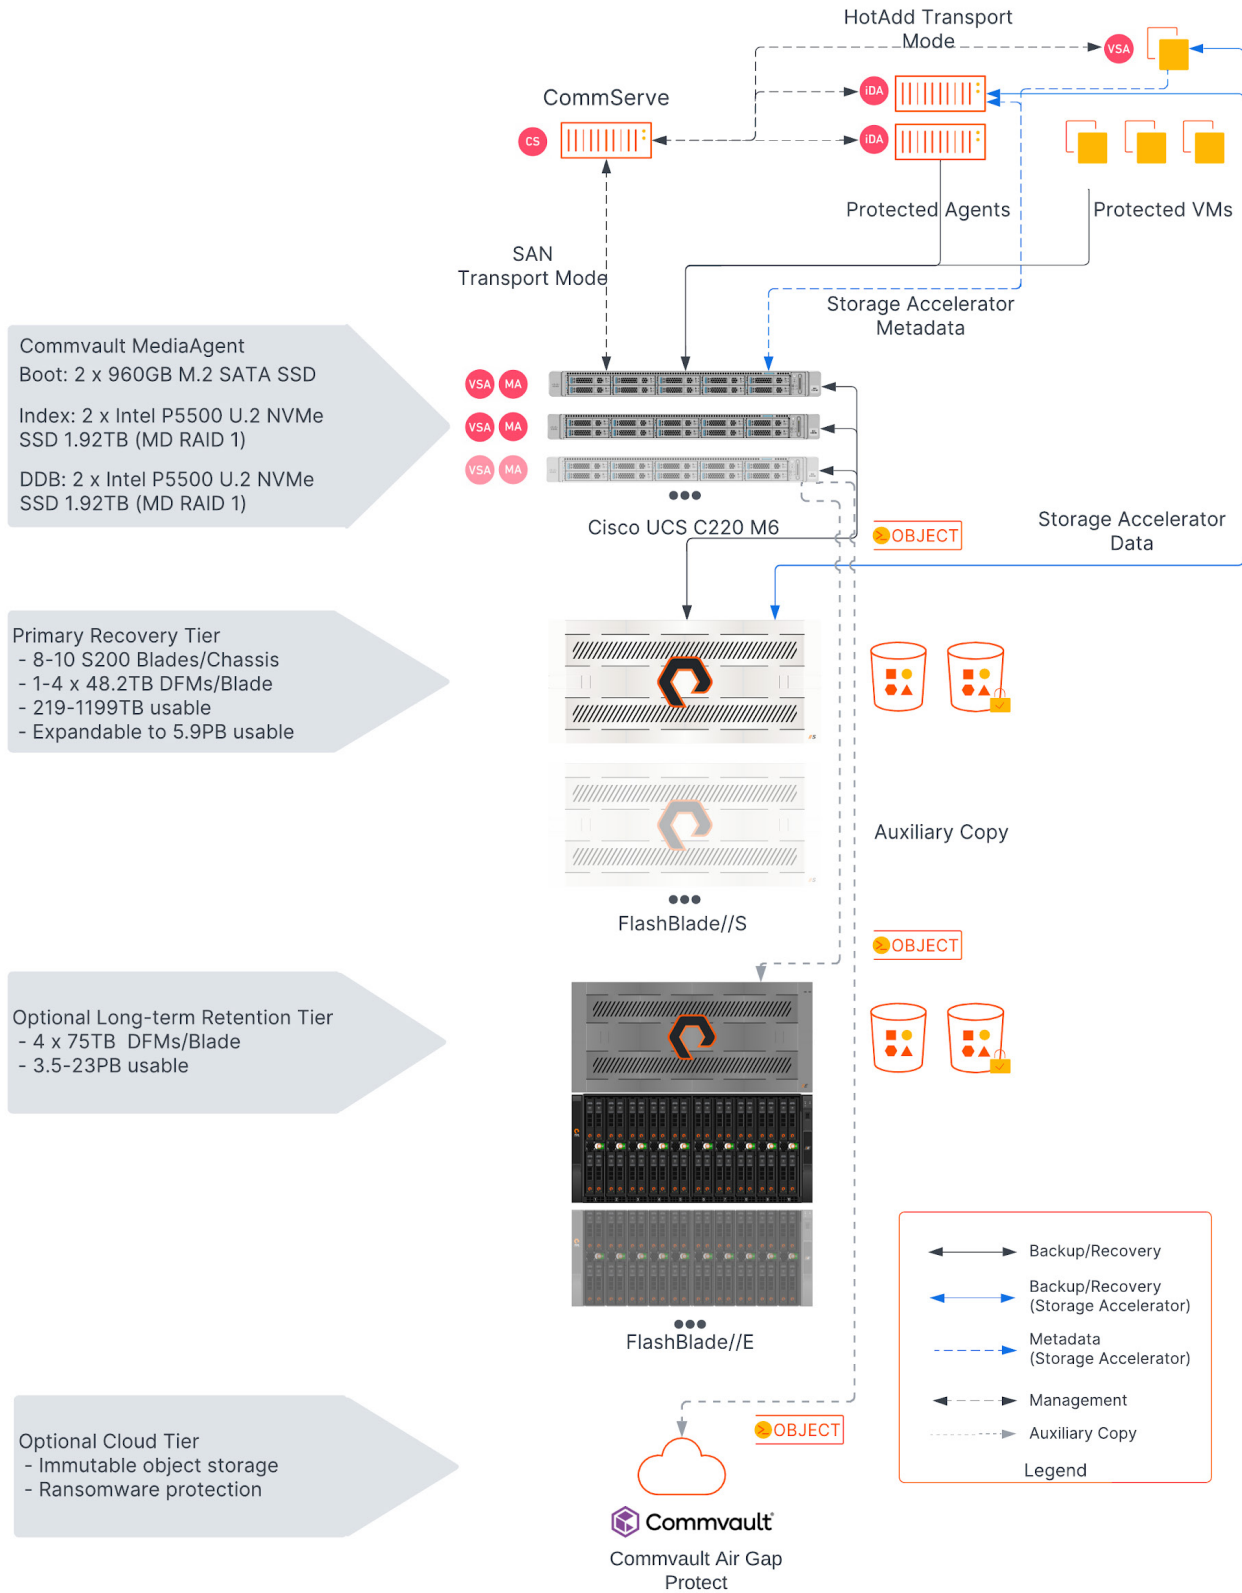


FIGURE 1 Logical architecture



## REFERENCE ARCHITECTURE

Figure 2 shows the physical architecture. Each component needs to be connected to two network switches for redundancy. The FlashBlade//S needs at least four switch ports, at least two from each switch, grouped in a link aggregation group (LAG). Grouping more ports will increase the available bandwidth.

If you wish to use optional fibre channel connectivity, connect each server to two storage area network (SAN) fabrics for redundancy.

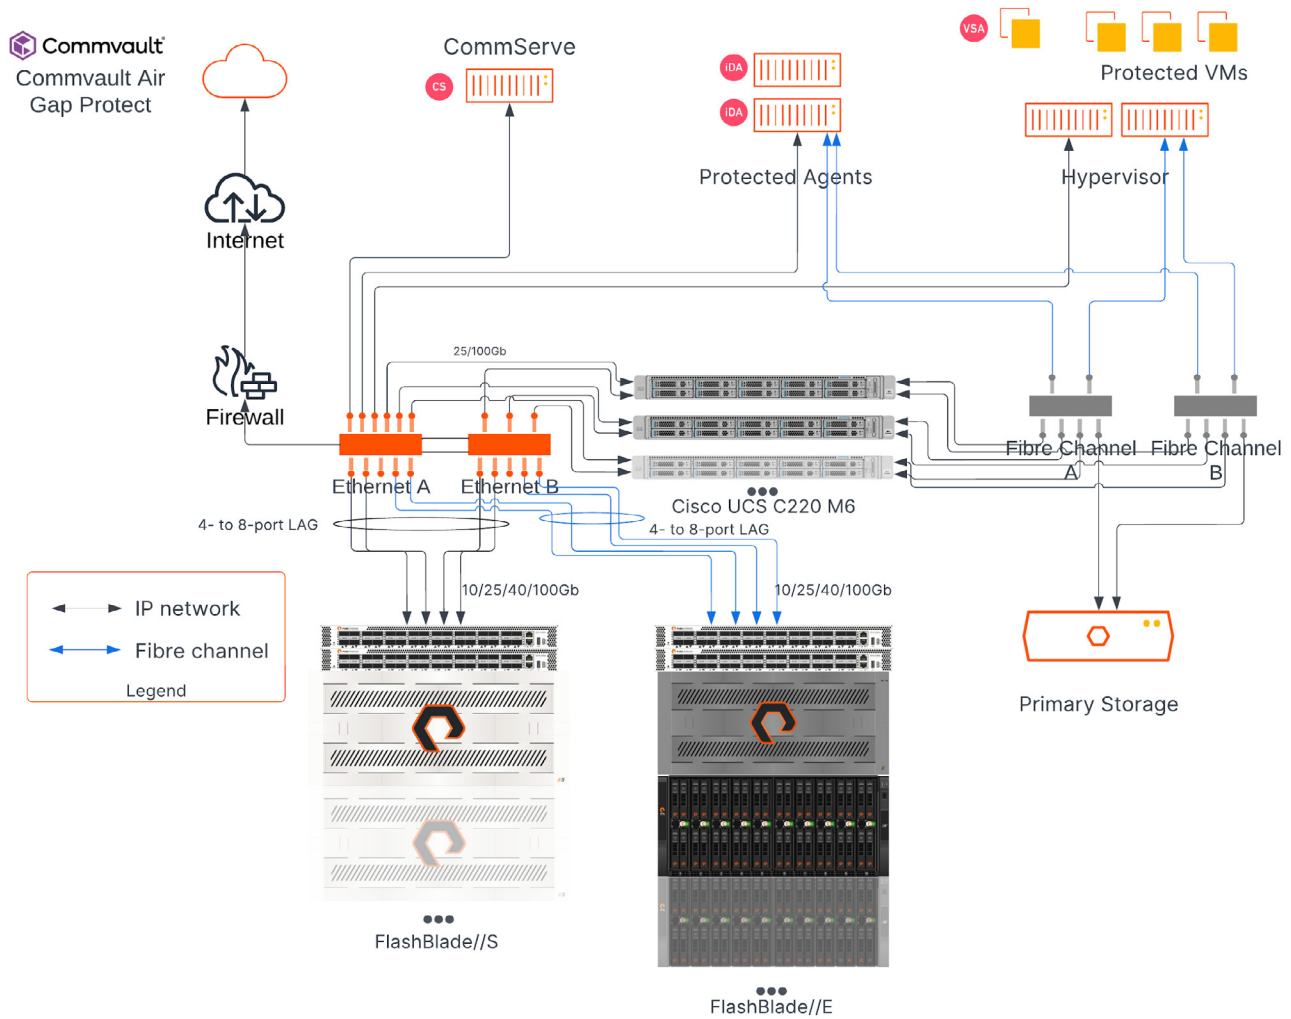


FIGURE 2 Physical architecture



## Core Solution Components

Commvault Cloud is an industry-leading platform for cyber resilience, delivering security, intelligence, and fast recovery for all your data. Commvault Cloud enables customers to protect, manage, and recover their data from one central cloud-based solution. Powered by Metallic AI, Commvault Cloud provides a layered defense, minimizing the impact of cyberattacks with early warning and cyber deception. Comprehensive threat scanning, remediation, intelligent quarantining, clean recovery validation, and unparalleled recovery speeds accelerate recovery. Commvault Cloud and Pure Storage FlashBlade make the world's fastest recovery predictable.

Cisco UCS C220 M6 servers offer high-performance data management in a small space. The C220 can process as much as half a petabyte per server but takes up only a single rack unit (RU). A base solution deployment includes two or more MediaAgent servers, and you can add servers to support as much storage as you need.

Pure Storage FlashBlade//S is the next generation of enterprise scale-out unified fast file and object (UFFO) storage, delivering rich data services with high density, capacity, performance, and scalability to meet the needs of modern applications. Using a distributed metadata architecture, FlashBlade//S offers multi-dimensional performance on a consolidated platform with NFS, SMB, and S3 protocol access. FlashBlade//S can scale seamlessly across multiple chassis to support up to ten petabytes of quad-level cell (QLC) flash storage. With an Evergreen Storage subscription, you can keep FlashBlade//S modern without data migrations.

## Customizable Solution Components

Protecting your backups from deletion and corruption is an important way to guard against cyberattacks. Commvault and FlashBlade combine to add object lock capabilities to your data protection environment. You can combine locked and unlocked backups on the same FlashBlade array. You can learn more about Commvault and FlashBlade object lock in [Minimizing the Impact of Cyberattack Recovery](#).

Pure Storage FlashBlade//E balances capacity and performance to deliver high density at an acquisition cost comparable to disk-based storage. Total cost of ownership (TCO) is significantly lower than comparable offerings thanks to industry-leading power efficiency and reliability. It's an ideal platform for data that has to be kept around for a long time and doesn't have to meet the tightest recovery service-level agreements (SLAs). Servicing the higher capacity may require additional MediaAgents.

Commvault Air Gap Protect gives organizations a way to easily meet data isolation requirements. As a cloud-based solution, it includes tamper-proof backups with multi-layered zero-trust access controls, rapid deployment, simplified management, and turn-key scalability. It further enhances your security posture with threat detection, including anomaly monitoring and honey pots.

Commvault IntelliSnap can orchestrate storage-based snapshots as part of a data protection workflow for a number of platforms and applications, and the MediaAgents can act as proxies to leverage those snapshots to offload backups away from production systems. If you wish to take advantage of this capability, and you wish to use fibre channel, your MediaAgents must include fibre channel adapters. The server specifications include the available options and recommendations.



## Designing for Your Environment

This section covers the elements required to design a deployment for your environment. Start by choosing the base configuration, then add optional components based on your needs.

This reference design does not specify networking equipment due to the number of available options and customer preferences. You must have a compatible network available to connect the solution components.

### Single-chassis Base Configurations

The base configurations can support from 195TB to 1.1PB of backup data on FlashBlade//S, or 90TB to 500TB of object locked data, after data reduction. You can scale the entire system to encompass 5.9PB usable storage on a single FlashBlade//S cluster. You will need one UCS C220 M6 server for each 500TB of backup data, with a minimum of two MediaAgents to provide resilience against server failure.

Table 1 shows the recommended starting configurations for deploying the solution with a single FlashBlade//S chassis. These FlashBlade//S models strike a balance between performance and storage, with roughly equal capacity intervals at each level. All options are equally expandable as your data grows. The MediaAgent server counts align with the storage sizing based on [Commvault deduplication building blocks](#).

Baseline Size	FlashBlade//S Usable Capacity (Object Lock effective capacity)	MediaAgent Servers	FlashBlade//S200 Model
Small	197TB (90TB)	2	1C8B1D-48.2 197TB
Medium	394TB (180TB)	2	1C8B2D-48.2 394TB
Large	699TB (318TB)	3	1C9B3D-48.2 699TB
Extra Large	1101TB (500TB)	4	1C10B4D-48.2 1101TB

TABLE 1 Solution baseline configurations

For more information on expanding the environment, refer to the [Expansion Guidelines](#) section.

## Commvault Software

The solution enables multiple use cases that leverage Commvault features. Table 2 details the minimum Commvault releases for these features.

Use Cases	Product	Commvault Release
Backup and recovery IntelliSnap Object lock	Commvault Cloud Foundational Protection or higher package	2023E or newer
Cloud-based isolated storage	Optional: Air Gap Protect	2023E or newer
Autonomous infection scanning	Optional: Threat Scan	2023E or newer
Cloud-based isolated recovery	Optional: Clean Room Recovery	2024 or newer

TABLE 2 Minimum Commvault versions by use case





### Pure Storage FlashBlade//S

The solution leverages FlashBlade//S200D arrays, with S200 blades and 48TB DirectFlash® Modules (DFMs). Base configurations use single-chassis arrays that can be expanded across multiple chassis. Multi-chassis systems require external fabric modules (XFM). FlashBlade//S arrays must be running Purity//FB release 4.1.9 or later for full compatibility; however, the latest 4.3.x release is recommended.

### MediaAgent Server Specifications

The MediaAgents in this reference architecture run a minimal installation of [Rocky Linux 8](#). This Red Hat-compatible distribution provides a stable foundation for Commvault's software and other components. [CIS Benchmark](#) Level 1 and SELinux add additional layers of security hardening to reduce the MediaAgent attack surface.

Table 3 lists the general MediaAgent hardware specifications. While this guide focuses on Cisco hardware, you can apply these specifications to other server products.

Item	Description	Quantity
<b>Core components</b>		
CPU	Intel Xeon Silver 4314 Processor 16 cores	2 (32 cores total)
RAM	Registered ECC DDR4 RDIMMs	512GB
Boot drive	SATA M.2 SSD 500GB minimum capacity RAID-1 configuration	2
Index drive	NVMe SSD, high performance, medium endurance 2TB-class	2
Deduplication database drive	NVMe SSD, high performance, medium endurance 2TB-class	2
<b>Customizable components</b>		
Networking	25Gbps or higher Ethernet ports 100Gbps preferred	2 ports minimum 4 ports recommended
Fibre channel (optional) for IntelliSnap backup copy and VMware SAN transport mode	16Gb or higher fibre channel host bus adapter (HBA) ports 32Gb HBA recommended	1-2 (2 total ports)

TABLE 3 MediaAgent hardware specifications



## Customizable Components

### Fibre Channel Host Bus Adapters

To connect the MediaAgents to a fibre channel (FC) storage area network, you should include at least two FC HBA ports, connected to separate fabrics, for redundancy. FC HBAs allow you to use Commvault IntelliSnap to manage storage-based snapshots and to use VMware SAN transport mode to directly back up VMs from SAN storage systems. Cisco UCS virtual interface cards (VICs) provide virtual FC HBAs with fibre channel over Ethernet (FCoE) functionality that you can use for this purpose. However, you may require physical HBAs in your environment. In this case, you can install two single-port HBAs or one dual-port HBA. The HBAs must be capable of 16Gbps or higher data rates; 32Gbps is recommended.

### Pure Storage FlashBlade//E

FlashBlade//E is available for new deployments in capacities of up to 24.3PB usable (30PB raw), in increments of approximately 1200TB usable (1500TB raw), using 75TB DFMs. As of this writing, arrays under five full chassis can be upgraded to a maximum of 12174TB usable (15PB raw). The array must run Purity//FB release 4.3.4 or later.

### Cisco UCS C220 M6 Part Numbers

Table 4 lists the Cisco part numbers to purchase that meet the MediaAgent server specifications. You can find more details on the parts and available options in the [Cisco UCS C220 M6 SFF Rack Server Spec Sheet](#).

Part Number	Description	Quantity
<b>System chassis</b>		
UCSC-C220-M6N	1	Cisco UCS X210c Compute Node 2S Intel 3rd Gen CPU without CPU, memory, drive bays, drives, VIC adapter, or mezzanine adapters (ordered standalone)
UCSC-RIS2H-220M6	1	Half-height riser 2
UCSC-RIS3H-220M6	1	Half-height riser 3
UCSC-PSU1-1200W	2	1200W Titanium power supply for C-Series Servers
<b>CPU and memory</b>		
UCS-CPU-I4314	2	Intel Xeon Silver 4314 Processor 16 cores
UCS-MR-X32G2RW	16	32 GB RDIMM DRx4 3200 (8Gb)
UCS-DIMM-BLK	16	UCS DIMM Blank
<b>Networking</b>		
mLOM (choose one option)		
UCSC-M-V25-04	1	Cisco UCS VIC 1467 quad port 10/25G SFP28 mLOM
UCSC-M-V100-04	1	Cisco UCS VIC 1477 dual port 40/100G QSFP28 mLOM
UCSC-M-V5Q50G	1	Cisco UCS VIC 15428 quad port 10/25/50G MLOM
UCSC-M-V5D200G	1	Cisco UCS VIC 15238 dual port 40/100/200G MLOM



Part Number	Description	Quantity
<b>Expansion (optional with 40/100G or 40/100/200G networking only, choose up to one option)</b>		
UCSC-PCIE-C100-04	1	Cisco UCS VIC 1495 Dual Port 40/100G QSFP28 CNA PCIe
UCSC-P-V5D200G-D	1	Cisco UCS VIC 15235 Dual Port 40/100/200G CNA PCIe
<b>Storage</b>		
UCS-M2-HWRAID	1	Cisco Boot optimized M.2 RAID controller (holds up to two M.2 SATA SSDs)
UCS-M2-960GB	2	960 GB M.2 SATA SSD
UCS-NVMEI4-I1920	4	1.9TB 2.5in U.2 Intel P5500 NVMe High Perf Medium Endurance
<b>Additional parts</b>		
Optional fibre channel HBA		
UCSC-P-Q6D32GF UCSC-P-B7D32GF	1	Cisco-QLogic QLE2772 2x32GFC Gen 6 Enhanced PCIe HBA Cisco-Emulex LPe35002-M2-2x32GFC Gen 7 PCIe HBA
<b>Transceivers and Optics</b>		
Cisco supports a number of transceiver and optics options, too many to adequately cover. You can review the supported options in the VIC data sheets. <a href="#">Cisco UCS Virtual Interface Card 1400/14000 Series Data Sheet</a> <a href="#">Cisco UCS Virtual Interface Card 15000 Series Data Sheet</a>		
<b>Power cables</b>		
Power cables vary by region and site <a href="#">Cisco UCS C220 M6 SFF Rack Server Spec Sheet</a>		
<b>Rack mounting hardware</b>		
Optional rail kit UCSC-RAIL-M6	1	Ball Bearing Rail Kit for C220 and C240 M6 rack servers
Optional cable management arm UCSC-CMA-C220M6	1	Reversible CMA for C220 M6 ball bearing rail kit

TABLE 4 Cisco UCS C200 M6 parts list



## Policy Considerations

The reference architecture gives you a wide variety of policy options. Commvault lets you define virtually limitless data lifecycles through storage pools and server backup plans. FlashBlade//S is capable of supporting multiple storage pools that have different retention and object locking settings. As you design your Commvault policies, there are several considerations to keep in mind:

- You should limit the number of storage pools and buckets you create. You do not need separate storage pools to handle different data types or retention, with the exception of object lock. A single storage pool can support multiple retention policies and backup agents.
- You should also limit the number of object lock retention options since each separate retention requires a separate storage pool. For optimal protection, you should deploy locked pools on more than one array in more than one site and use Auxiliary Copy to replicate data between them. Retention does not necessarily need to match between sites, so you can create a fan-in topology.
- You should use FlashBlade//S as the target for initial backups and your rapid recovery tier. Having an object locked storage pool as the first data copy gives your backups immediate protection against cyberattack. Use a separate pool for longer-term retention needs.
- For long-term retention and archive, you should minimize data churn in the deduplication stores. Selective copy–replicating backup data from only some of your backups–reduces the amount of volatile data that the deduplication database (DDB) has to process, which in turn reduces the MediaAgent load and storage requirements. FlashBlade//E should be your target for on-premises retention. Commvault Air Gap Protect can serve cloud-based needs and enables Commvault's [Clean Room Recovery](#) for ransomware recovery in the cloud.

## System Placement and Site Design

This section will assist with designing appropriate system placement in single- and multi-site environments.

### General Guidelines

For the best backup and recovery performance, you should place the MediaAgents and FlashBlade//S in the same site as the data you're protecting. This avoids bottlenecks from crossing wide-area network (WAN) links. However, the solution can support backups from remote sites. Commvault's network throttling allows you to limit bandwidth utilization, and source-side deduplication and compression reduce WAN traffic and backup times.

You should minimize network hops between MediaAgents and FlashBlade by assigning them IP addresses on the same IP subnet as each other.

### Multiple Sites

If you are deploying the solution in two or more sites, the main consideration for site design is recovery needs for each site. If you need the same recovery speed and capacity in each site, deploy identical configurations. If your performance requirements are lower in a secondary site, you may, for example, choose to deploy only FlashBlade//E and not FlashBlade//S.

Regardless of your storage decisions, you should deploy redundant MediaAgents in each site where you will place a FlashBlade//S. While a MediaAgent can manage storage in multiple locations, having local MediaAgents already available gives you better resilience after a site disaster.

Commvault policy design needs to account for multiple sites. You can define any progression of data between storage targets to meet your unique needs.



Figure 3 illustrates a two-site deployment. Backups land on FlashBlade//S in Site 1. Auxiliary Copy replicates the data to FlashBlade//S in Site 2. Auxiliary Copy within each site selectively replicates long-term data from FlashBlade//S to FlashBlade//E and from Site 2 to Commvault Air Gap Protect.

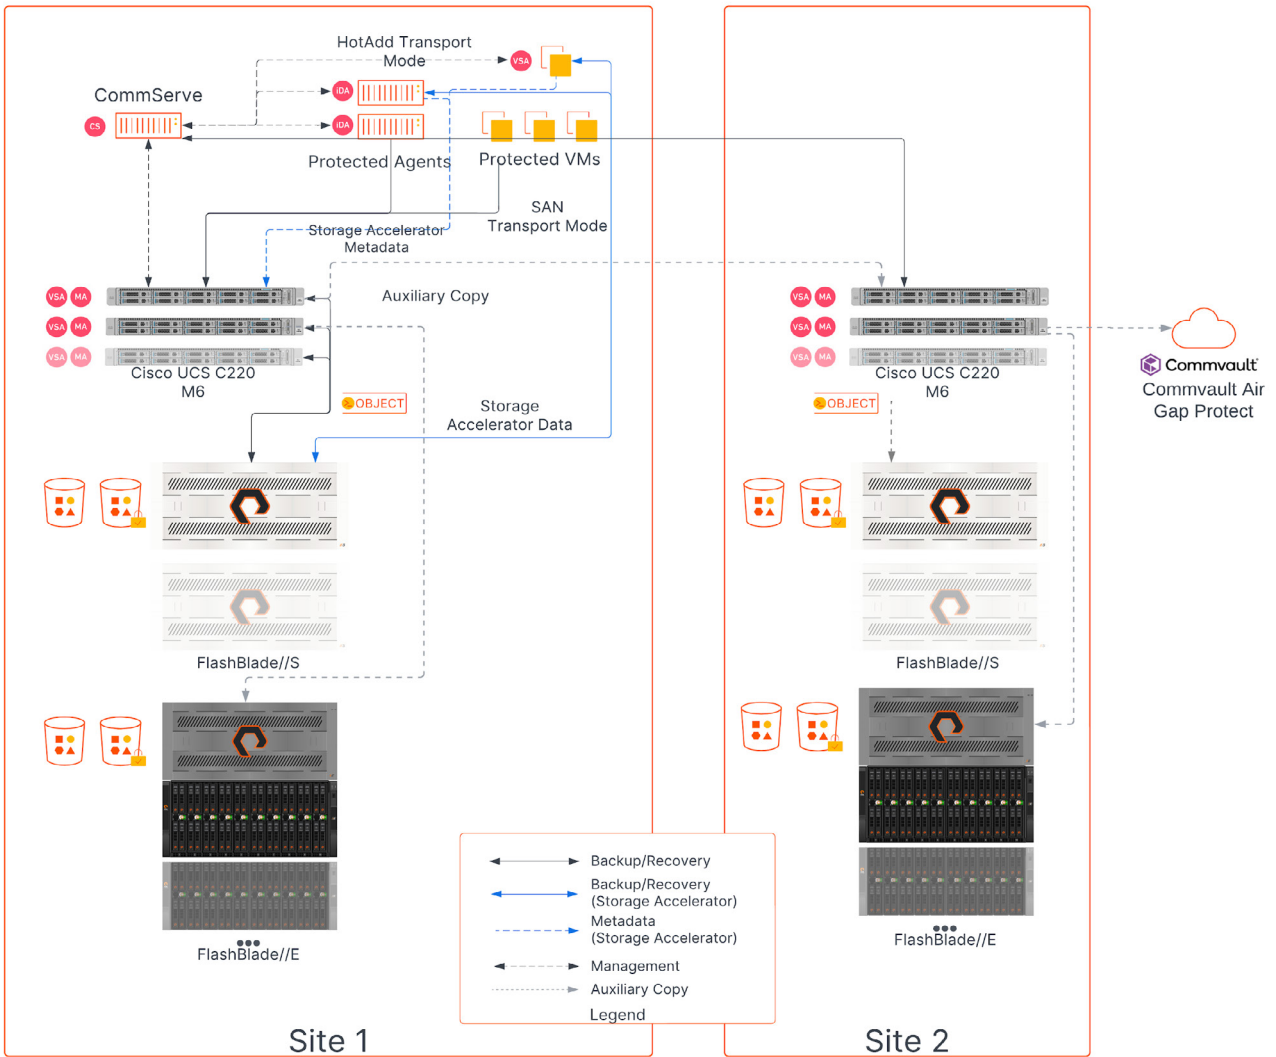


FIGURE 3 Two-site deployment

### Inter-site Networking

Ensure that the network links between your sites have enough available bandwidth to meet your SLAs. Commvault will reduce bandwidth requirements by only sending new unique data between sites, and it can throttle network traffic between sites to prevent saturation. Recovery SLAs are especially important to these calculations; site links are typically the narrowest part of the data path, so undersized links will prevent you from meeting your targets.

## Expansion Guidelines

As your environment grows, you will reach points where you need to expand your backup and recovery solution to increase capacity, performance, or both. This section offers guidance on what increases to make and when.

### Predicting Future Needs

Pure1®, the Pure Storage cloud-based management portal, includes capacity analysis and planning tools you can use to align expected growth and budget cycles. You should periodically review your predicted growth to stay ahead of future capacity needs.

### What to Add

In most environments, capacity will be the main driver for expanding both FlashBlade footprint and number of MediaAgents. You will need to add blades or DFMs, and eventually chassis, to FlashBlade//S, and Commvault will need additional MediaAgents at certain capacity points.

You may find that you are not facing capacity pressure but need more parallelism in backup or recovery. In that case, you can add MediaAgents without changing the FlashBlade//S.

### FlashBlade//S Capacity Tiers

If you need less than 1100TB of FlashBlade//S storage, we recommend upgrading from your initial configuration to one of the configurations in Table 5.

FlashBlade//S Usable Capacity (Object Lock effective capacity)	Minimum MediaAgent Servers	FlashBlade//S200 Model
394TB (180TB)	2	1C8B2D-48.2 394TB
699TB (318TB)	3	1C9B3D-48.2 699TB
1101TB (500TB)	4	1C10B4D-48.2 1101TB

TABLE 5 FlashBlade//S Capacity Tiers

Exceeding 1100TB requires adding one or more chassis to the cluster, which requires additional rack space but is a seamless operation. You should plan to add capacity in sets of two blades with four 48TB DFMs each, or around 220TB usable capacity, at a time. You should also plan to add a MediaAgent along with each capacity increase.

### FlashBlade//E Capacity Tiers

As of this writing, FlashBlade//E with 75TB DFMs allows upgrades in increments of 1500TB raw flash, or approximately 1200TB usable, to a maximum of 12TB usable in five chassis. Expansion beyond five chassis will be supported in a future release.

If you are using the FlashBlade//E to store only long-term data and not all backup jobs, you should not need additional MediaAgents. The MA requirements increase with the size and frequency of data copies, and copying every backup to FlashBlade//E would require the same number of MAs per PB as for FlashBlade//S.

## Performance

The architecture is designed to make full use of the available performance of the FlashBlade. However, there may be situations where external factors like client bandwidth hold you back. In that case, you may need additional MediaAgents to enable more parallel streams. You will also need additional storage pools to take advantage of the MAs.



## Conclusion

Reducing IT risk and cost starts with simplifying. Simpler solution design makes acquisition easier and cheaper. Validation of a prescribed Cisco UCS configuration leads to better, more consistent outcomes with less investment.

Pure and Commvault achieve exactly this to serve your data protection and cyber resilience needs. A bill of materials takes out any guesswork from acquisition. A standardized configuration deploys easily and scales and expands seamlessly as your organization grows.

With Pure Storage, you can do away with backup storage downtime and migrations. Non-disruptive upgrades and expansion have changed the way Pure customers manage maintenance. Evergreen Storage options let you keep arrays modern without stealing dollars and resources from strategic projects every few years.

Options for additional storage tiers let you optimize for cost, with long-term retention on premises with FlashBlade//E or in the cloud with Commvault Air Gap Protect. Both options offer the same type of protection from cyberattacks as FlashBlade//S and let you maintain simplicity without sacrificing performance.

When you're ready to learn more, visit Pure Storage [Commvault solutions page](#), and contact your Pure Storage account team or reseller.

## Additional Resources

### Next Steps

- Learn more about [FlashBlade//S](#) and [FlashBlade//E](#).
- See how [Commvault Cloud](#) can modernize data protection in your environment.
- Discover how [Cisco UCS C-Series Rack Servers](#) can simplify your infrastructure management.

### Supporting Information

- [FlashBlade//S Data Sheet](#)
- [FlashBlade//E Data Sheet](#)
- [Pure1](#)
- [Commvault Cloud](#)
- [Commvault Air Gap Protect](#)
- [Commvault Clean Room Recovery](#)
- [Reference Architecture Deployment Guide](#)

## About the Author



Roy Child is a Senior Solution Architect with Pure Storage, responsible for defining data recovery solutions and reference architectures for primary workloads such as Oracle, SQL, and VMware. Roy has worked in and with the data protection industry for over 20 years, from end user to IT architect with multiple backup and recovery products, followed by product management with Commvault. Roy joined Pure Storage in April, 2019.