White Paper
Public

Cisco Hyperflex for Video Surveillance Management with Cohesity and Milestone XProtect

Contents

Executive summary

Today, physical safety and security are mandatory for organizations throughout the world. A critical component of physical safety and security is video surveillance and analytics associated with it. Video surveillance is no longer a simple stream of video data that is stored and retrieved only after a certain incident happens to do forensic analysis. With the current world situation, real-time alarms and threat detection are highly desirable to prevent any unwanted incidents. These features are especially important in the context of smart city solutions, to make them smarter and more efficient.

Data collection is the first step in a video surveillance solution. Planners and developers can then use and act on this data to improve security. However, connectivity is never guaranteed, and therefore data must be safeguarded at all times. One of the key challenges to realize the full potential of video surveillance and predictive analytics is the traditional infrastructure used to store and manage data. Traditional infrastructures are siloed in nature due to inherent scalability limitations. These siloed environments are expensive and create complex data management challenges at scale and inhibit enterprises to move towards modern approaches which could help reduce cost while providing better ROI.

To reduce costs and simplify the lifecycle of video surveillance data from inception to insight, Cisco, Milestone, and Cohesity have developed a next-generation architecture based on Cisco HyperFlex™ hyperconverged systems, Milestone XProtect Corporate video management software (VMS), and Cohesity DataPlatform software on the Cisco Unified Computing System™ (Cisco UCS®) for simplified data management and essentially limitless scalability. This video surveillance solution allows enterprises to:

- Centrally manage and store video
- Support large storage capacities
- Use cost-effective S3 storage in the cloud for long-term data retention
- Scale essentially limitlessly to accommodate future growth
- Record high frame rates without dropping frames
- Handle high-resolution video

Cisco HyperFlex unifies compute, storage, networking from the core to the edge and provides best-in-class Hyperconverged Infrastructure, while Milestone XProtect Corporate provides an industry-leading video management software, and Cohesity DataPlatform provides a single, software-defined data management platform from edge to core to cloud.

Solution overview

In this architecture, Cisco HyperFlex, a hyperconverged system is used to host Milestone video management software (VMS) system in a virtual environment. Virtualization can make better use of hardware resources, and XProtect Corporate's powerful distributed server architecture is well suited to virtualization and scaling to keep up with an organization's growth. One of the main benefits of virtualization is that a virtualized deployment allows organizations to install Milestone VMS components individually and provides high availability for all the components.

The operating system disk and data disk for live database storage is carved out of the Cisco Hyperflex datastore. Cisco HyperFlex Data Platform stripes and replicates data across all nodes in the cluster. Data remains available if one or more components fail (depending on the replication factor used). If a virtual

machine is migrated to a new location or if any component fails, the data platform does not require data movement because any virtual machine can read data from any location. Thus, the movement of virtual machines has no performance impact or cost.

In addition, the Cisco HyperFlex HX Data Platform makes data instantly accessible and highly available, with always-on deduplication and compression that reduces storage needs.

Cisco HyperFlex systems

The Cisco HyperFlex is engineered on the Cisco UCS platform. It is part of a data center architecture that supports traditional, converged, and hyperconverged systems with common policies and infrastructure management.

Cisco HyperFlex solutions enable enterprises to deploy applications from the core to the edge to multicloud environments. These solutions power mission-critical applications and databases in the data center and enable multicloud development and deployment of cloud-native applications. These solutions extend the simplicity of hyperconvergence to the edge at a distributed scale.

Cisco HyperFlex systems with Intel® Xeon® Scalable processors deliver hyperconvergence with power and simplicity for any application, on any cloud, anywhere. Engineered on Cisco UCS, Cisco HyperFlex systems deliver the agility, scalability, and pay-as-you-grow economics of the cloud with the benefits of multisite, distributed computing at global scale.

- **Any application**: Cisco HyperFlex systems support virtualized and containerized deployment with multiple hypervisors and have been tested and validated for numerous enterprise applications.
- **Any cloud**: The platform includes tools for application performance monitoring, application placement, and cloud mobility so that you can deploy your applications and place them wherever your business needs dictate.
- **Anywhere**: You can achieve true global scale with the simplicity of hyperconverged infrastructure that reaches to your network edge. Cloud-based deployment, management, and monitoring scales through templates that make deploying hundreds of remote sites as simple as deploying only one.

Cisco HyperFlex systems combine these features:

- **Software-defined computing** in the form of nodes based on Cisco UCS servers
- **Software-defined storage** with the powerful Cisco HyperFlex HX Data Platform software
- **Software-defined networking** with Cisco UCS fabric that integrates smoothly with the Cisco® Application Centric Infrastructure (Cisco ACI™) solution
- **Cloud-based management** with Cisco Intersight™ software as a service and multicloud container support from Cisco Container Platform for Cisco HyperFlex systems

Cisco Intersight platform

Cisco Intersight is a management platform delivered as a service with embedded analytics for your Cisco and third-party IT infrastructure. This platform offers an intelligent level of management that enables IT organizations to analyze, simplify, and automate their environments in more advanced ways than the prior generations of tools. Cisco Intersight platform provides an integrated and intuitive management experience for resources in the traditional data center and at the edge. With flexible deployment options to address complex security needs, getting started with the Cisco Intersight platform is quick and easy.

The Cisco Intersight platform has deep integration with Cisco UCS and with Cisco HyperFlex systems, allowing remote deployment, configuration, and ongoing maintenance. The model-based deployment works for a single system in a remote location or hundreds of systems in a data center, and enables rapid, standardized configuration and deployment. It also simplifies the maintenance of those systems whether you are working with small or large configurations.

Cisco Intersight software delivers a new level of cloud-powered intelligence that supports lifecycle management with continuous improvement. It is tightly integrated with the Cisco Connected Technical Assistance Center (TAC) and Smart Call Home. Expertise and information flow seamlessly between the Cisco Intersight, Cisco UCS and Cisco HyperFlex users. Remediation and problem resolution are supported with automated uploading of error logs for rapid root-cause analysis.

Cisco Intersight platform uses a Continuous Integration (CI) and Continuous Deployment (CD) approach, so you never have to worry about whether your software is up-to-date. Even the Cisco Intersight virtual appliance is automatically updated.

Cisco Intersight software offers these main features:

- Software-as-a-service (SaaS)-based management: SaaS delivers global management with frequent updates that don't impede your operations.
- Proactive guidance: The recommendation engine provides notifications, insights, and actionable intelligence to ease daily operations.
- Security and extensibility: The service is designed for secure connection and data access with an extensible architecture for third-party integrations.
- Enhanced support: Enhanced capabilities and Cisco TAC integration help you quickly respond to problems before they affect operations.
- Intuitive experience: Help your administrators and DevOps teams be more effective, less burdened with details, and more productive.

Figures 1 and 2 provide an overview the Cisco Intersight platform.

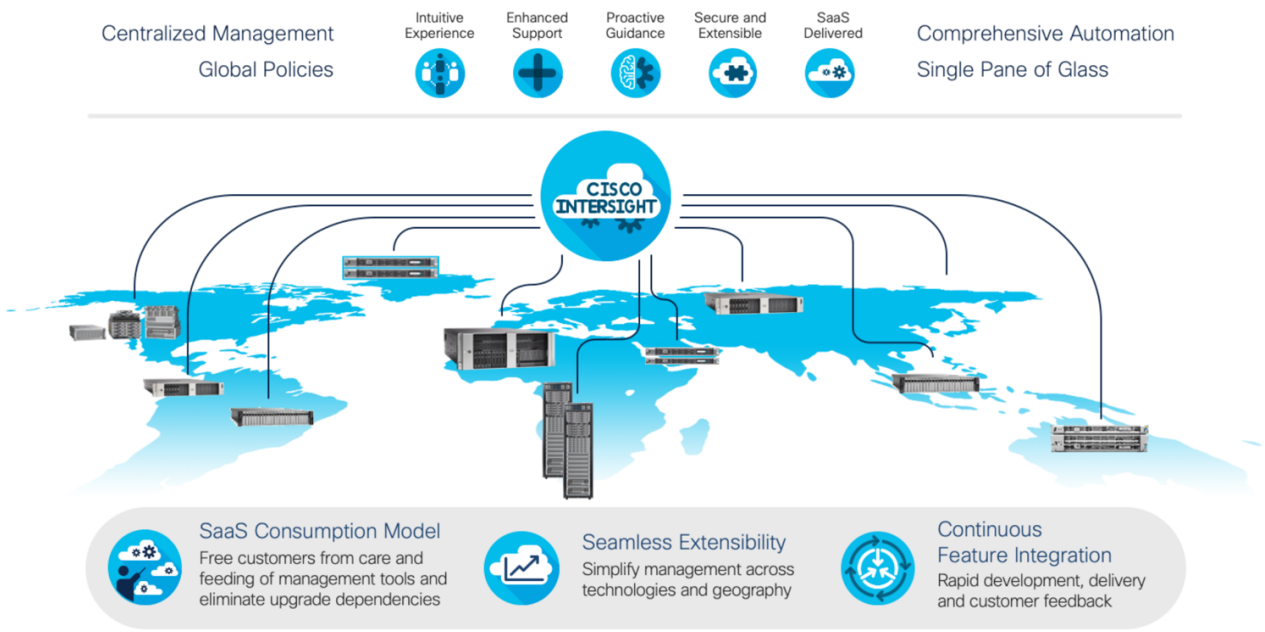

Figure 1.
Cisco Intersight cloud-based infrastructure automation

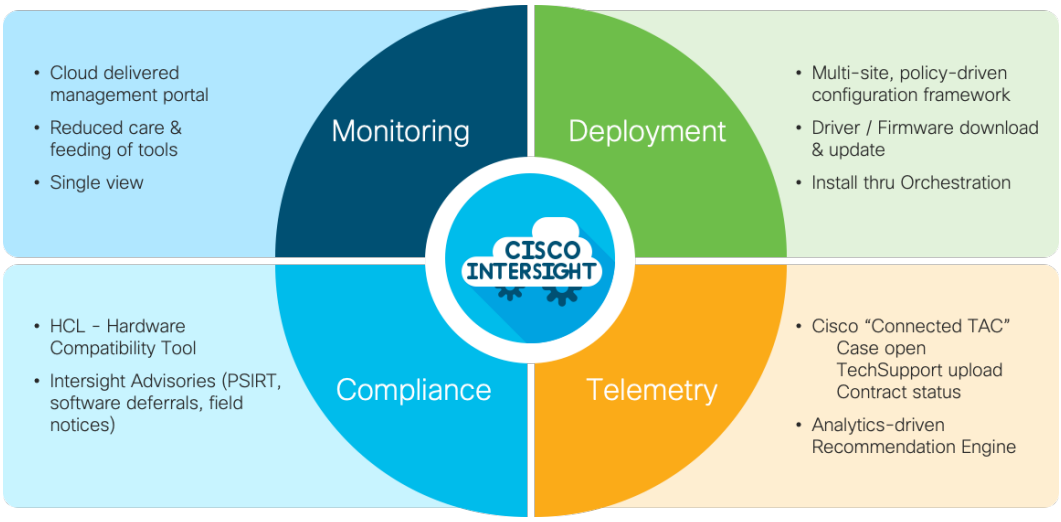

Figure 2.
Cisco Intersight core capabilities: Infrastructure operations

One of the biggest challenges organizations face is managing data center servers while extending services wherever data or users reside with massive scale. The Cisco Intersight platform provides a single interface that can access hundreds of centralized and remote locations and provides support for the entire infrastructure lifecycle.

The Cisco Intersight platform is developed, integrated, and tested using the Cisco Secure Development Lifecycle guidelines. This secure product development and deployment practice has several components, including inherent design and development practices, implementation testing, and creation of a set of recommendations for deploying with the highest levels of security. Cisco development processes are ISO 27001 certified, with a certification specific to Cisco Intersight development currently in the audit pipeline.

Please see the Cisco Intersight help pages for details about managing roles and resources.

Cohesity DataPlatform: Redefining Data Management

The vast majority of enterprise data—backups, archives, file shares, object stores, and data used for development and testing and analytics—sits in fragmented infrastructure silos, making the data hard to protect, expensive to manage, and difficult to analyze. Cohesity on Cisco UCS consolidates silos onto one web-scale platform that spans on-premises, cloud, and edge environments and uniquely empowers organizations to run applications on that platform so they can more easily back up and extract insights from their data. Certified with a Cisco Validated Design, the joint solution complements the Cisco HyperFlex platform and is available directly from Cisco, simplifying purchasing, deployment, and support.

Cohesity SpanFS file system

At the core of the Cohesity DataPlatform is a fully distributed, shared-nothing file system. Inspired by web-scale principles, Cohesity SpanFS is a unique file system that is meticulously designed to address the challenge of mass data fragmentation.

To effectively consolidate data, enterprises need a file system that can handle the requirements of multiple use cases simultaneously. To meet modern data management requirements, SpanFS provides the following features (Figure 3):

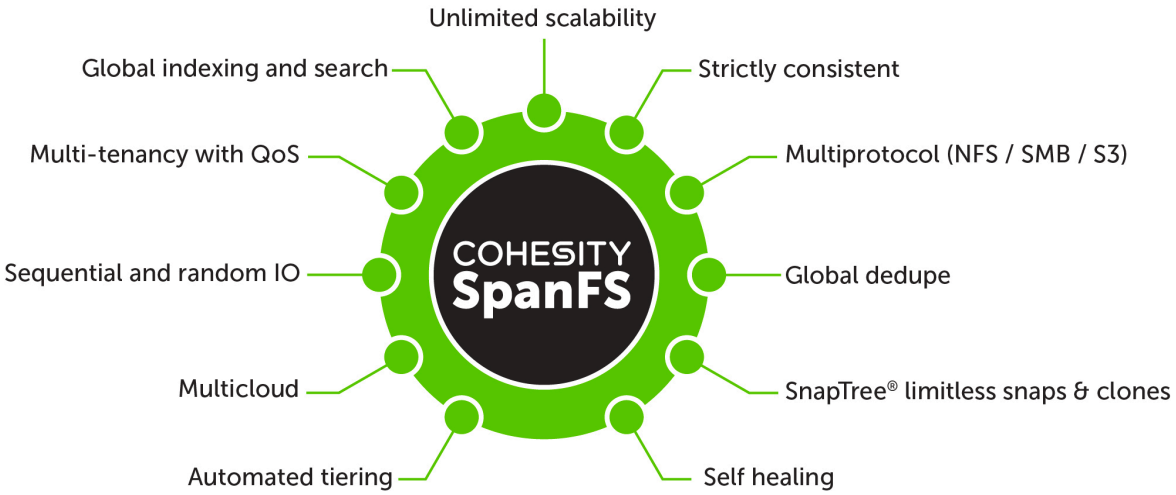

Figure 3.
Cohesity SpanFS features

- Unlimited scalability: Start with as few as three nodes and grow limitlessly on-premises or in the cloud with a pay-as-you-grow model.
- Strict consistency: Help ensure data resiliency with strict consistency across nodes in a cluster.
- Multiprotocol support: Support traditional Network File System (NFS) and Server Message Block (SMB)-based applications as well as modern S3-based applications. Read and write to the same data volume with simultaneous multiprotocol access.
- Global deduplication: Significantly reduce your data footprint by deduplicating across data sources and workloads with global variable-length deduplication.
- Unlimited snapshots and clones: Create and store an unlimited number of snapshots and clones with significant space savings and no performance impact.

- Self-healing design: Auto-balance and auto-distribute workloads across a distributed architecture.
- Automated tiering: Automatic data tiering across solid-state disk (SSD), hard-disk drive (HDD), and cloud storage helps you achieve the right balance between cost optimization and performance.
- Multicloud support: Natively integrate with leading public cloud providers for archiving, tiering, and replication and to protect cloud-native applications.
- Sequential and random I/O: Achieve high I/O performance by auto-detecting the I/O profile and placing data on the most appropriate media.
- Multitenancy with quality of service (QoS): Natively support multiple tenants with QoS, data isolation, separate encryption keys, and role-based access control (RBAC).
- Global indexing and search: Rapidly perform global searches as a result of file and object metadata indexing.

The Cohesity DataPlatform file system provides full at-rest encryption based on the strong 256-bit Advanced Encryption Standard (AES-256). The Cohesity DataPlatform architecture provides this strong data security while maintaining the flexibility to use the available hardware and software resources. Cohesity DataPlatform also supports multifactor authentication (MFA).

Milestone Systems

Milestone Systems is a global leader in open-platform IP video surveillance software. Milestone has provided easy-to-use, powerful video management software in more than 200,000 installations worldwide.

Milestone XProtect provides open-architecture products that are compatible with more IP cameras, encoders, and digital video recorders than products from any other manufacturer. Because Milestone provides an open platform, you can integrate today's best business solutions and expand your capabilities with future innovations. Visit www.milestonesys.com for more information.

XProtect Corporate

XProtect Corporate is a powerful IP VMS solution designed for large-scale and high-security deployments. Its single management interface enables efficient administration of the system, including all cameras and security devices, regardless of the system's size or whether it is distributed across multiple sites. For systems demanding supreme situational awareness and precise response to incidents, XProtect Corporate includes Milestone XProtect Smart Wall. XProtect Corporate includes advanced video grooming functions and encryption capabilities that help organizations reduce video storage costs while helping ensure the integrity of video evidence and compliance with industry and federal regulations.

VMS server components

The video management software includes a number of server components

Management server

The management server is the central component of the VMS system. It stores the configuration of the surveillance system in a Microsoft SQL database, either on a Microsoft SQL Server on the management server computer itself or on a separate SQL Server on the network. It also handles user authentication, user rights, the rule system, and more. To improve system performance, you can run several management servers as a Milestone Federated Architecture. The management server runs as a service and typically is installed on a dedicated server.

Users connect to the management server for initial authentication, then transparently to the recording servers for access to for video recordings, etc.

Recording server

The recording server is responsible for communicating with the network cameras and video encoders, recording the retrieved audio and video, and providing client access to both live and recorded audio and video. The recording server is also responsible for communicating with other Milestone products connected through the Milestone Interconnect technology.

Event server

The event server handles various tasks related to events, alarms, maps, and third-party integrations through the Milestone Integration Platform (MIP) software development kit (SDK).

- Events
  - All system events are consolidated in the event server so that partners have one place and interface for implementing integrations that use system events.
  - The event server also offers third-party access to the generic events or analytics events interface to send events to the system.
- Alarms
  - The event server hosts the alarm feature, alarm logic, alarm state, and alarm database. The alarm database is stored in the same SQL database that the management server uses.
- Maps
  - The event server hosts the maps that are configured and used in the XProtect Smart Client MIP SDK.
  - Third-party-developed plug-ins can be installed on the event server and have access to system events.

Log server

The log server stores all log messages for the entire system in a SQL database. This log messages SQL database can reside on the same SQL Server as the management server's system configuration SQL database or on a separate SQL Server. The log server typically is installed on the same server as the management server, but it can be installed on a separate server for increased performance of the management and log servers.

Microsoft SQL Servers and databases

The management server, event server, and log server store, for example, the system configuration, alarms, events, and log messages in SQL databases on one or more SQL Server installations. The management server and the event server share the same SQL database, and the log server has its own SQL database. The system installer includes Microsoft SQL Server Express, which is a free edition of Microsoft SQL Server.

For very large systems or systems with many transactions to and from the SQL databases, Milestone recommends that you use the Microsoft SQL Server Standard or Microsoft SQL Server Enterprise edition of SQL Server on a dedicated computer on the network and on a dedicated hard disk drive that is not used for other purposes. Installing the SQL Server on its own drive improves the performance of the entire system.

Mobile server

XProtect Mobile server handles logins to the system from XProtect Mobile client or XProtect Web Client.

A XProtect Mobile server distributes video streams from recording servers to XProtect Mobile client or XProtect Web Client. This offers a secure setup where recording servers are never connected to the Internet. When a XProtect Mobile server receives video streams from recording servers, it also handles the complex conversion of codecs and formats allowing streaming of video on the mobile device.

Client components

XProtect Management client

The management client is a single-point system administration interface with robust features for configuration and day-to-day management of the system. It provides centralized management of all aspects of system configuration.

XProtect Smart Client

Designed for Milestone XProtect IP video management software, XProtect Smart Client is an easy-to-use client application that provides intuitive control over security installations. Manage security installations with XProtect Smart Client to give users access to live and recorded video, instant control of cameras and connected security devices, and an overview of recordings.

XProtect Web Client

XProtect Web Client is a client designed for the occasional or remote user that needs easy access to live monitoring, playback and export. XProtect Web Client also provides access to activating system events and outputs.

XProtect Mobile client

The XProtect Mobile client is a client designed for the user on the go. It offers easy access to live monitoring, playback and export of video, as well as access to activating system events and outputs.

You can use the XProtect Mobile client as a remote recording device by using the device's built-in camera and the Milestone Video Push feature. With Video Push activated, video from the device's camera is streamed back to the VMS and recorded as if it is a standard camera.

Solution architecture

The solution architecture includes four Cisco HyperFlex HX240c M5L Nodes for hosting Milestone components, infrastructure virtual machines and storing live recording database, and three Cisco UCS S3260 Storage Servers for deploying Cohesity DataPlatform to store the archive database.

The Cisco HyperFlex system is used primarily for live database recording, with a retention time of 3 days, and Cohesity DataPlatform software-define storage (SDS) is used for archiving. The retention time can be modified according to the storage configuration.

Physical topology

Cisco UCS and Cisco HyperFlex systems are managed and configured through the Cisco Intersight, a lifecycle management platform for the infrastructure. The Cisco Intersight platform includes the Cisco HyperFlex installer in all editions, providing an easy way to deploy Cisco HyperFlex clusters.

All you need to do to deploy Cisco HyperFlex clusters is to connect power and network cables and claim the servers in the user interface. Apply a cluster profile to a Cisco HyperFlex cluster through the Cisco Intersight platform, and your systems or clusters are configured automatically in minutes.

Cisco Intersight supports managed Cisco UCS S-Series servers with limited features and functionalities, and the complete management features will be made available in the future. The Cisco UCS S-series servers are configured through Cisco UCS Manager and the Cohesity DataPlatform software installation is done through ISO image file, which will automate the process of installing the underlying Linux operating system, copy the Cohesity software packages, and prepare the nodes for the initial setup of the Cohesity cluster.

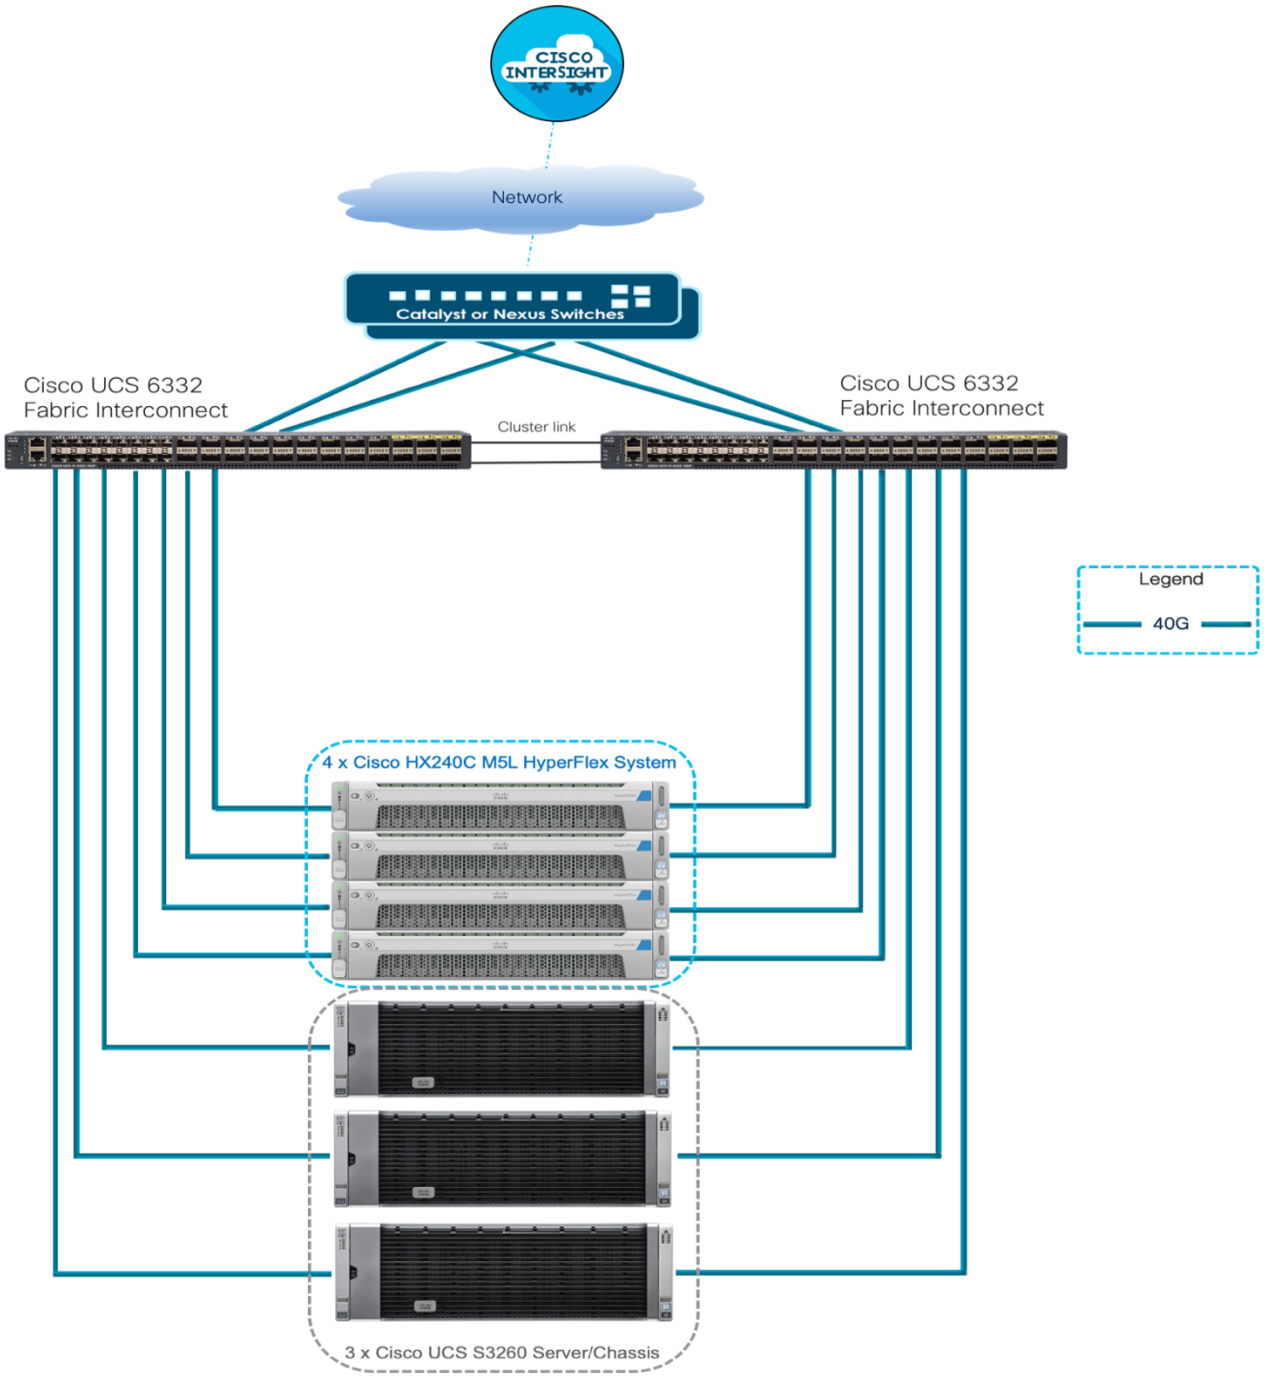

Figure 4.
Physical topology

The Milestone VMS system has several components that handle specific tasks, and its distributed architecture allows the system to scale. All the VMS components and infrastructure virtual machines are deployed on the Cisco HyperFlex system, and the live video streams are stored directly in high-performance, highly scalable Cisco HyperFlex storage. The four-node Cisco HyperFlex hybrid system is configured with a replication factor of 3 (RF3) and provides a usable storage capacity of 55.3 TB, which can be easily scaled according to your requirements. Cohesity DataPlatform configured on three Cisco

UCS S3260 servers provides a total raw capacity of 587 TB, which also can be easily scaled. The live database recording server configuration with Cisco HyperFlex storage was used for Milestone validation and certification purposes only. The actual configuration can vary according to the environment and performance requirements.

Tables 1 and 2 show the hardware components and software versions used in the solution.

Table 1. **Table 1** Hardware components

Hardware components			
Component	**Model**	**Quantity**	**Comments**
Cohesity storage node	Cisco UCS S3260 M5 Chassis	3	1 x UCS S3260 M5 Server Nodes per Chassis (Total = 3nodes) Per Server Node: 2 x Intel Xeon Gold 6240 (2.6GHz/18 cores), 256 GB RAM Cisco UCS S3260 Dual Pass Through Controller based on Broadcom IT Firmware Cisco UCS S3260 System IO Controller with 1300-series VIC included UCS S3260 240G Boot SSD (Micron 6G SATA) SSD for OS 4 x Cisco UCS S3260 Top Load 3X 3.2 TB SSD 21 x Cisco UCS S3260 10TB (512e) Top Load Intel i350T4 quad-port 1G copper with iSCSI NIC
Cisco HX hyperconverged infrastructure	Cisco HX240C M5L HyperFlex System	4	Each Server Configuration: 2 x Intel Xeon Gold 6248 (2.5GHz/20cores), 384 GB RAM 1 x 240GB M.2 6G SATA SSD for ESXi Hypervisor 1 x 240GB 2.5 inch Enterprise Value 6G SATA SSD for System / Logs 1 x 3.2TB Enterprise performance 12G SAS SSD(3X endurance) for Cache 6 x 8TB 12G SAS 7.2K RPM LFF HDD (4K) for Data/Capacity 1 x Cisco 12G Modular SAS HBA
UCS UCS Fabric Interconnects	Cisco UCS 6332 Fabric Interconnects	2	

Table 2. Software versions

Software distributions and versions		
Layer	**Component**	**Version or Release**
Storage (Chassis) UCS S3260	Chassis Management Controller	4.0(2k)
	Shared Adapter	4.3(2k)
Compute (Server Nodes) UCS S3260 M5	BIOS	S3260M5.4.0.2k
	CIMC Controller	4.0(4k)
Cisco HyperFlex system: Cisco HyperFlex HX240c M5L	BIOS	C240M5.4.0.4j
	CIMC Controller	4.0(4e)
Cisco UCS 6332-16UP Fabric Interconnect	UCS Manager	4.0(4e)
	Kernel	5.0(3)N2(4.02c)
	System	5.0(3)N2(4.02c)
Software		
Cisco HyperFlex		4.0(1b)
ESXi hypervisor		6.7 U2
Windows Server Standard Edition		2019
Milestone XProtect Corporate		2020R1b

Logical topology

Figure 5 provides a logical view of the solution architecture. Milestone VMS components and infrastructure virtual machines are deployed on the Cisco HyperFlex system with recording servers directly writing live recording database on Cisco HyperFlex storage. Cohesity DataPlatform is used to store the archived database data, which can then be moved to the public cloud on S3 storage for long-term retention (LTR). Having the recording servers on a virtual machine helps you scale the system when new cameras are added. As more cameras are added, more virtual machines are added, with no need to configure storage or networking components.

With the Cisco HyperFlex hyperconverged system, storage and computing resources can be scaled independent of each other. The topology in Figure 5 shows 15 recording server virtual machines, each recording 100 camera streams, thus capturing a total of 1500 camera streams. Three of the VMware ESXi hosts are configured with four recording server virtual machines, and one of the ESXi hosts is configured with three recording server virtual machines. Infrastructure virtual machines such as the Milestone management server, SQL Server, and XProtect client can be deployed individually as virtual machines and in this solution, all these Milestone components are deployed in a single virtual machine for certification purposes.

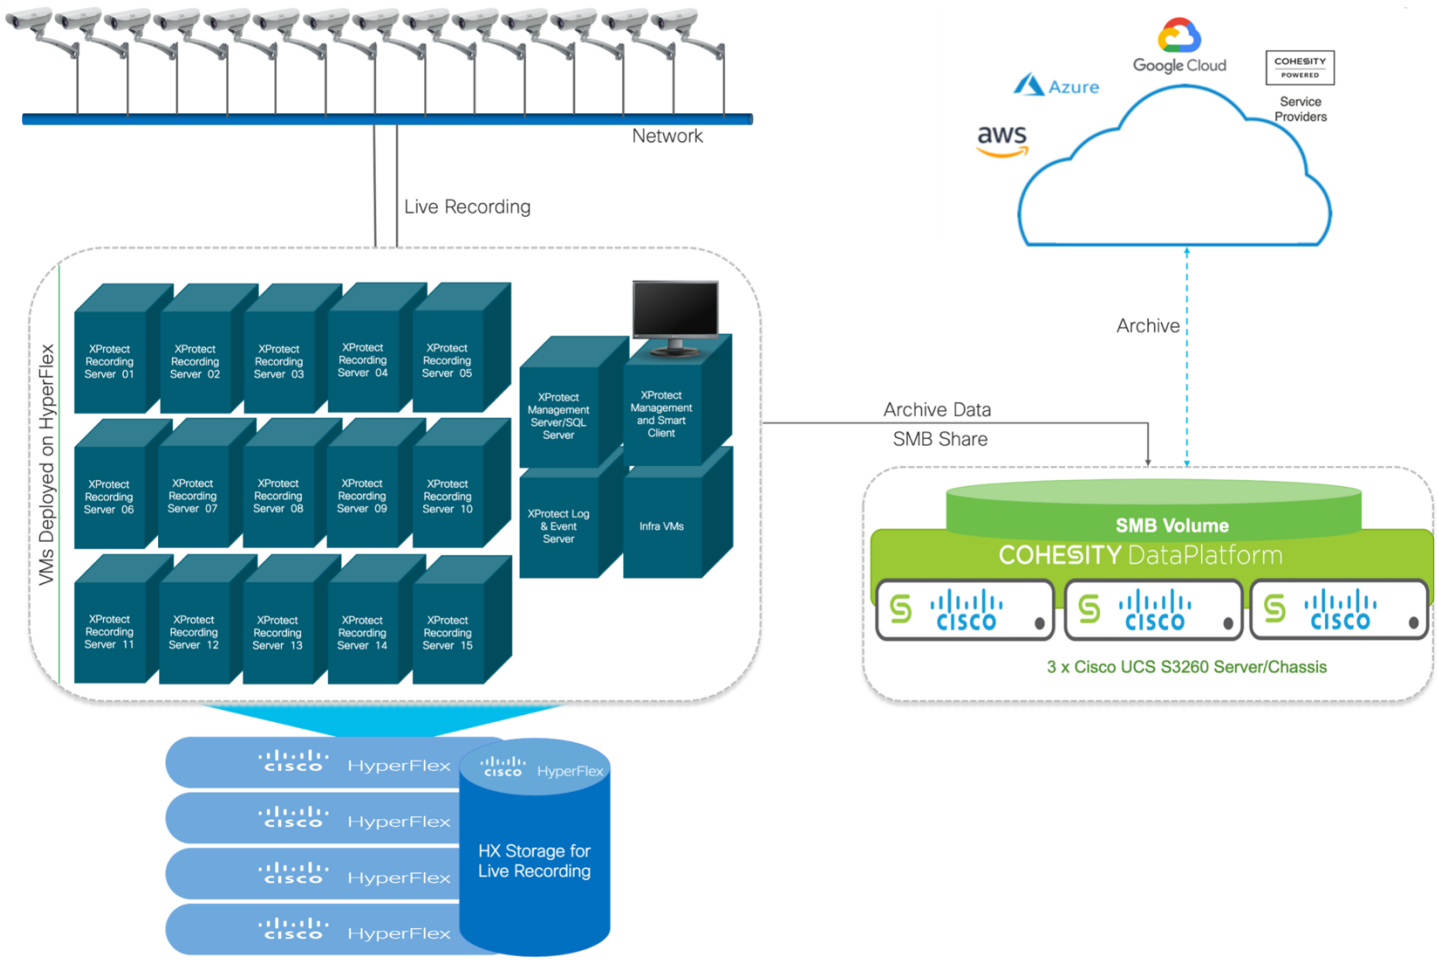

Figure 5.
Logical topology

Table 3 summarizes the XProtect components and infrastructure virtual machine configuration details. Each recording server is configured with 10 vCPUs, 32 GB of RAM, a 100-GB OS disk, 5 TB of space for live video recording, and 100 camera devices.

Table 3. Milestone virtual machine configuration details

VM Name	Quantity	vCPU	RAM	Datastore	
				OS disk	Data disk
XProtect Management server, Management Client, SQL Express, log & event server	1	10	32 GB	400 GB	
XProtect Recording server	15	10	32 GB	100 GB	5 TB per VM
Feed server	2	4	8 GB	100 GB	
XProtect Smart Client	1	8	16 GB	100 GB	

All the video streams were configured to use 1920 x 1080 full high-definition (HD) H.264 codec video with the number of frames set to 30 per second. Fifteen recording servers, with each recording server receiving video feed from 100 cameras, were simulated. A total of 1500 camera feeds were recorded on the system, with the retention time set to 3 days. The data was then archived on Cohesity network-attached storage (NAS) for 1 week and then copied to the Amazon Web Services (AWS) cloud on S3 storage for 3 months.

Note: For systems with more than 100 cameras, Milestone recommends that you use dedicated servers for all or some of the components.

Note: By default, Microsoft Windows formats a New Technology File System (NTFS) disk with a 4-KB block size. For best performance, make sure that your storage is formatted with an NTFS 64-KB block size, regardless of the use of RAID and JBOD.

Archiving

Milestone's unique multistage storage technology allows you optionally to use internal and external video archives. This capability offers the possibility of using cost-effective, high-density storage systems such as scale-out NAS for long-term video storage. To further reduce the cost of long-term video storage, multistage video storage offers the possibility of reducing the frame rate at each archiving step. Using this grooming capability, you can significantly reduce the cost of storage when operating with long video retention times.

Cohesity DataPlatform provides globally distributed NFS, SMB, and S3 object storage with best-in-class global deduplication and compression. Volumes are provisioned as Views on Cohesity. Cohesity Views provide network-accessible storage that is distributed across the Cohesity cluster using NFS volumes, SMB or Common Internet File System (CIFS) mount paths, or S3 compliant object-based storage. The SMB protocol is enabled on the Cohesity View and mounted as an SMB file share on the recording servers to archive the data.

With global deduplication and compression and its next-generation data management capabilities, Cohesity supports further cost reduction by natively using S3 object storage on the premises or in a multicloud environment for long-term retention (archiving) based on data governance policies and processes across the organization.

Note: For an even performance profile across all components in the data path, it is recommended to stream archive data to Cohesity in staggered fashion from each of the recording servers. This approach prevents the recording servers from all writing archive database data sets to Cohesity at the same time. It also helps ensure that ingress performance of the camera feeds on the recording servers is not affected by heavy read I/O processing during archiving.

Network design

The connectivity of the solution is based on 40-Gbps connectivity, and traffic is isolated using VLANs. Figure 6 shows the network layout of the Cisco HyperFlex system and Cohesity.

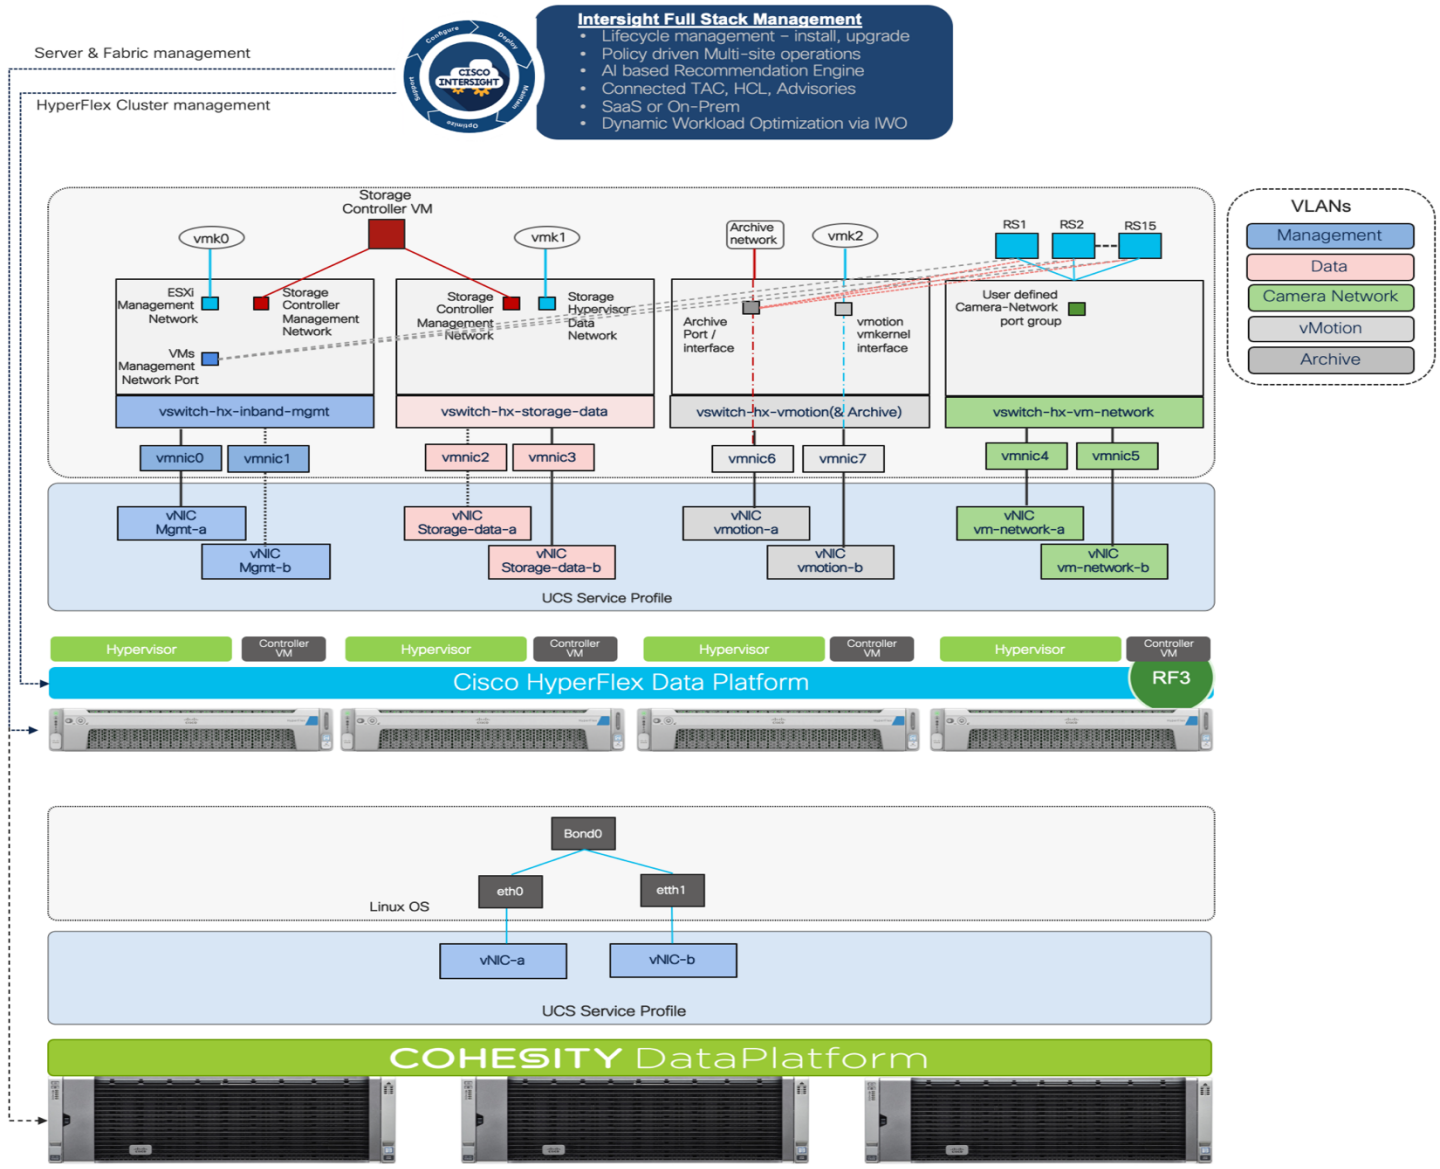

Figure 6.
Logical network layout

Notes:

- Dotted lines represent a standby link.
- All “a” vNICs connect to FI-A ,and all “b” vNICs connect to FI-B.
- A maximum transmission unit (MTU) 0f 9000 is needed for Storage-data, vMotion, Archive DB, and Camera-Networks.

Cisco HyperFlex network settings

The Cisco HyperFlex HX Data Platform installer automatically creates service profiles, vSwitches, and VLANs based on user input. A total of four vNIC pairs (eight vNICs) are created; each pair has one vNIC from Fabric A and one from Fabric B.

Each ESXi host needs the following networks.

- Management traffic network: From VMware vCenter; handles hypervisor (ESXi server) management and storage cluster management
- Data traffic network: Handles the hypervisor and storage data traffic
- vMotion network: Used for virtual machine and storage vMotion traffic
- Archive network: Handles the archive database traffic
- Camera network: Handles the camera feed traffic

Four vSwitches, each with a pair of vNICs, are created for management, storage, vMotion (archive database), and Camera traffic, with each carrying a different network:

- vswitch-hx-inband-mgmt: This vSwitch is used for ESXi management and storage controller management
- vswitch-hx-storage-data: This vSwitch is used for ESXi storage data and HX Data Platform replication. These two vSwitches are further divided into two port groups with assigned static IP addresses to handle traffic between the storage cluster and the ESXi host.
- vswitch-hx-vmotion: This vSwitch is used for virtual machine, storage vMotion, and archive database traffic. The vMotion port group has a vNIC from Fabric B configured as active, and the archive database port group has a vNIC from Fabric A configured as active.
- vswitch-hx-vm-network: This vSwitch is used for camera feed traffic. You can add or remove VLANs on the corresponding vNIC templates in Cisco UCS Manager. See Managing VLANs in Cisco UCS Manager and Managing vNIC templates in Cisco UCS Manager for detailed steps. To create port groups on the vSwitch, see Adding Virtual Port Groups to VMware Standard vSwitch.

Table 4 shows the VLAN and vSwitch configuration used in the solution described in this document.

Table 4. VLAN and vSwitch configuration

VLAN type	Description	vSwitch	Active vNIC	Standby vNIC
VLAN for ESXi and HyperFlex Management Traffic	VLAN and port group name: hx-inband-mgmt VLAN ID: 300	vswitch-hx-inband-mgmt	vmnic0	vmnic1
VLAN for HyperFlex storage traffic	VLAN and port group name: hx-storage-data VLAN ID: 400	vswitch-hx-storage-data	vmnic3	vmnic2
VLAN for VM vMotion	VLAN and port group name: hx-vmotion VLAN ID: 600	vswitch-hx-vmotion	vmnic7	vmnic6
VLAN for Archive data traffic	VLAN and port group name: archive VLAN ID: 200	vswitch-hx-vmotion	vmnic6	vmnic7
VLAN for camera feed network	VLAN and port group name: camera-network VLAN ID: 500	vswitch-hx-vm-network	vmnic4, vmnic5	--

Note: By default, the hx-vm-network vSwitch is configured as active-active. All other vSwitches are configured as active-standby.

The Cisco HyperFlex system can be integrated with the Cisco ACI solution, bringing software-defined networking (SDN) to the system, with innovations and capabilities that go far beyond what traditional SDN solutions provide. Cisco HyperFlex infrastructure, when connected to the Cisco ACI solution, extends the software-defined model into the data center network to deliver a scalable, application-centric, policy-based infrastructure for enterprise data centers.

The Cisco ACI solution can manage the virtual networking in a Cisco HyperFlex environment using a Cisco Application Policy Infrastructure Controller (APIC) controlled VMware vSphere Distributed Switch (vDS) or Cisco ACI Virtual Edge (AVE). The network can attach directly to virtual machines and physical servers with increased security, real-time monitoring and telemetry, and automated performance optimization.

Cohesity network settings

Two Cisco UCS vNICs are configured per node: one on the A-side fabric and one on the B-side fabric. The two interfaces are configured as slave interfaces in a bond in the Linux operating system, using bond mode 1 (active-passive).

A floating virtual IP address is assigned, one per node, and is used by Cohesity for all management, backup, and file services access. The address assignment is handled by the Cohesity software, and addresses are reassigned to an available node if any node should go offline. These floating addresses are all assigned in the Domain Name System (DNS) to a single A record, and the DNS server must respond to queries for that A record using DNS in a round-robin process.

Network configuration for Milestone validation

For validation and certification purposes, the network load was generated by the feed server, which is connected to a pair of Cisco Nexus switches. Figure 7 shows the virtual machine mappings to port groups and the IP address range used in the solution. The Milestone XProtect Management server, Management Client, infrastructure virtual machines, and all other XProtect component virtual machines are connected to the management port group. The recording servers are connected to the camera network, management and archive network port group and the feed server is connected to the camera network port group.

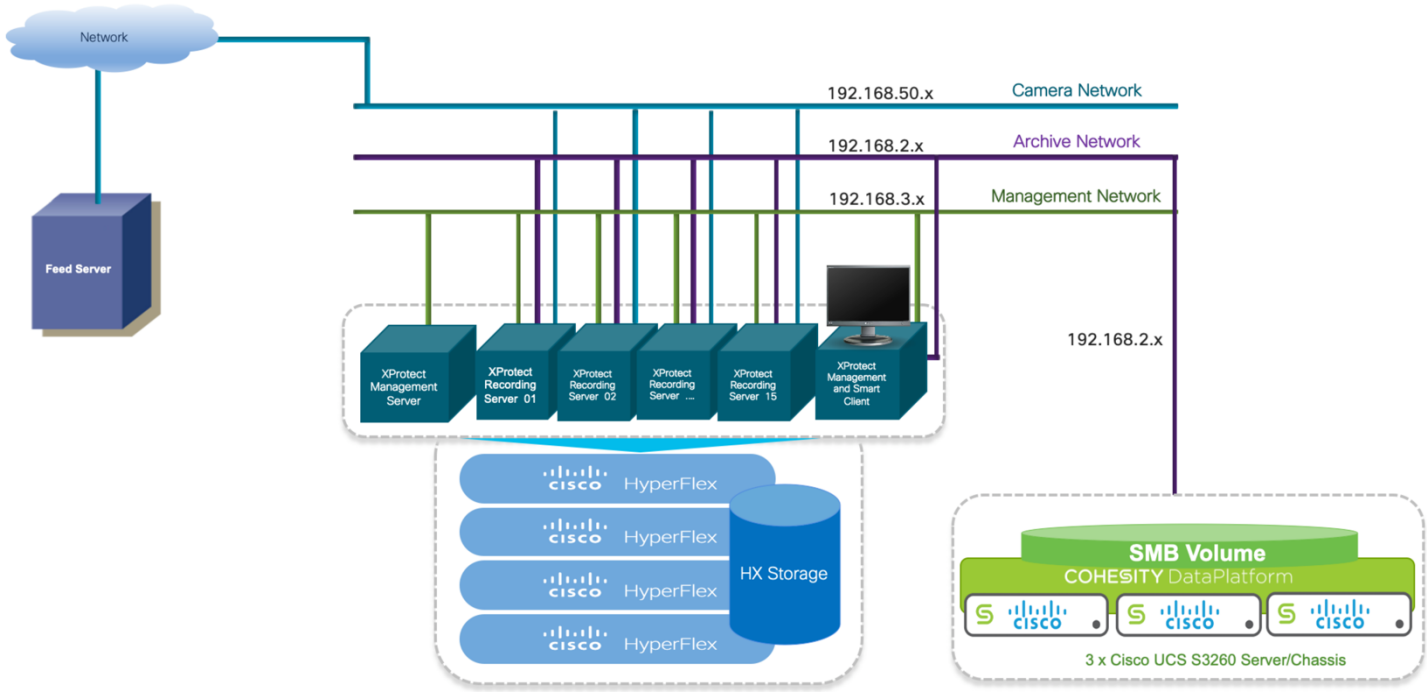

Figure 7.
Virtual machine and port group mappings

Cohesity configuration for archiving

Cohesity simplifies data management by consolidating silos onto one web-scale platform that spans on-premises, cloud, and edge environments and uniquely empowers organizations to run applications on that platform so organizations can more easily back up and extract insights from their data (Figure 8).

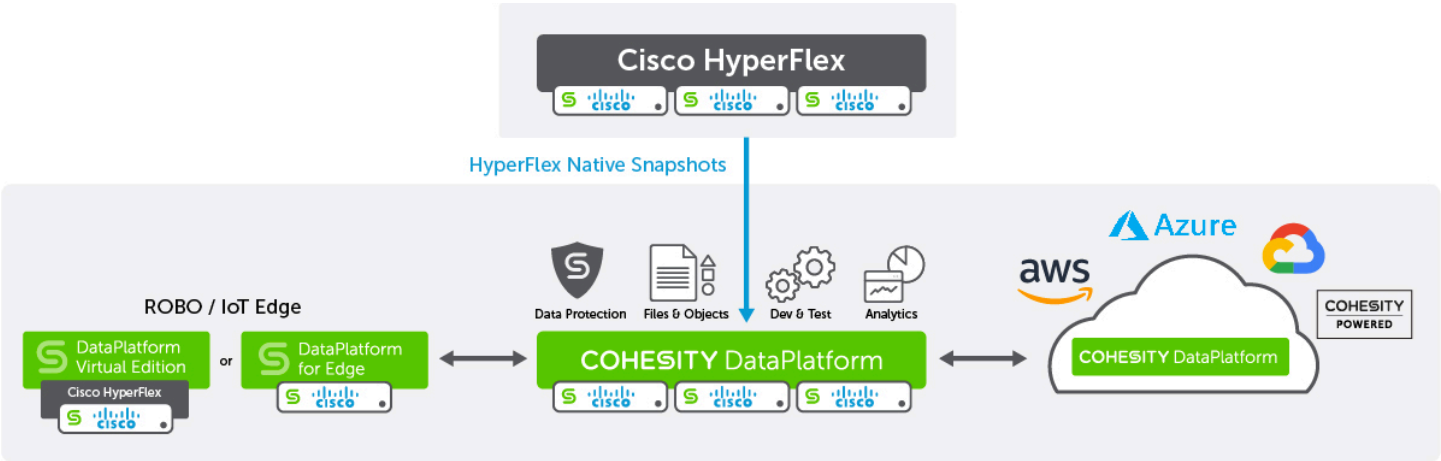

Figure 8.
Cohesity simplifies data management through a consolidated web-scale platform

The Cohesity SmartFiles feature (see additional information here) provides multiprotocol (SMB, NFS, and S3) NAS capabilities from Cohesity's patented SpanFS and SnapTree file systems. A Cohesity View provides a storage location with NFS, SMB, S3, and Swift mount paths in a storage domain in the Cohesity cluster.

For the validation reported in this document, the Cohesity SmartFiles feature was used to provide Milestone recording servers with a View and SMB share from Cohesity for the archive database workload. Cohesity provides a number of QoS policies that can be changed dynamically. QoS profiles have different settings based on different types of workloads; therefore, based on the storage I/O characteristics of the workload (for example, whether I/O is sequential or random), a suitable QoS policy can be used to meet performance requirements. The SMB share for Milestone was configured with journaled sequential dump QoS policy and with fast durable handles to boost performance. The fast durable handles option provides faster performance (more I/O operations per second [IOPS], metadata operations, file listing, etc.) for SMB clients. The feature can tolerate restarts of the Cohesity data service and short-duration network failures. However, it may not tolerate a node failure or SMB server service failure.

Configuring Cohesity Views

Cohesity Views represent mount points into a specific storage domain. The Views provide NFS or SMB and CIFS protocol access for files, snapshots, and clones of other views. QoS can be set for each View to tune performance for the target workload (Figure 9).

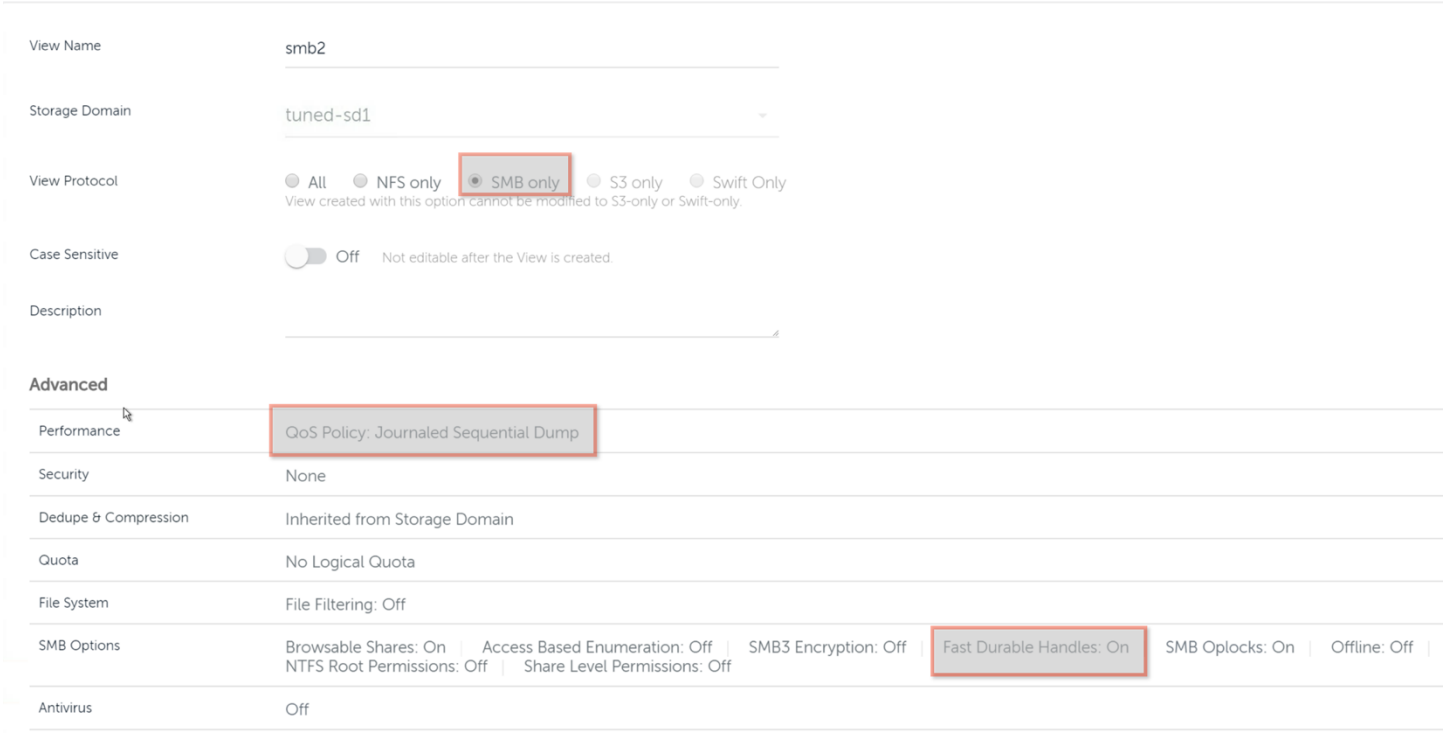

Figure 9.
Configuring a Cohesity View

Configuring protection policies

Protection policies describe backup frequency, backup retention, replication policies, and archive schedules. After policies are defined, they can be applied to protection jobs, ensuring appropriate protection strategies across your environment (Figure 10).

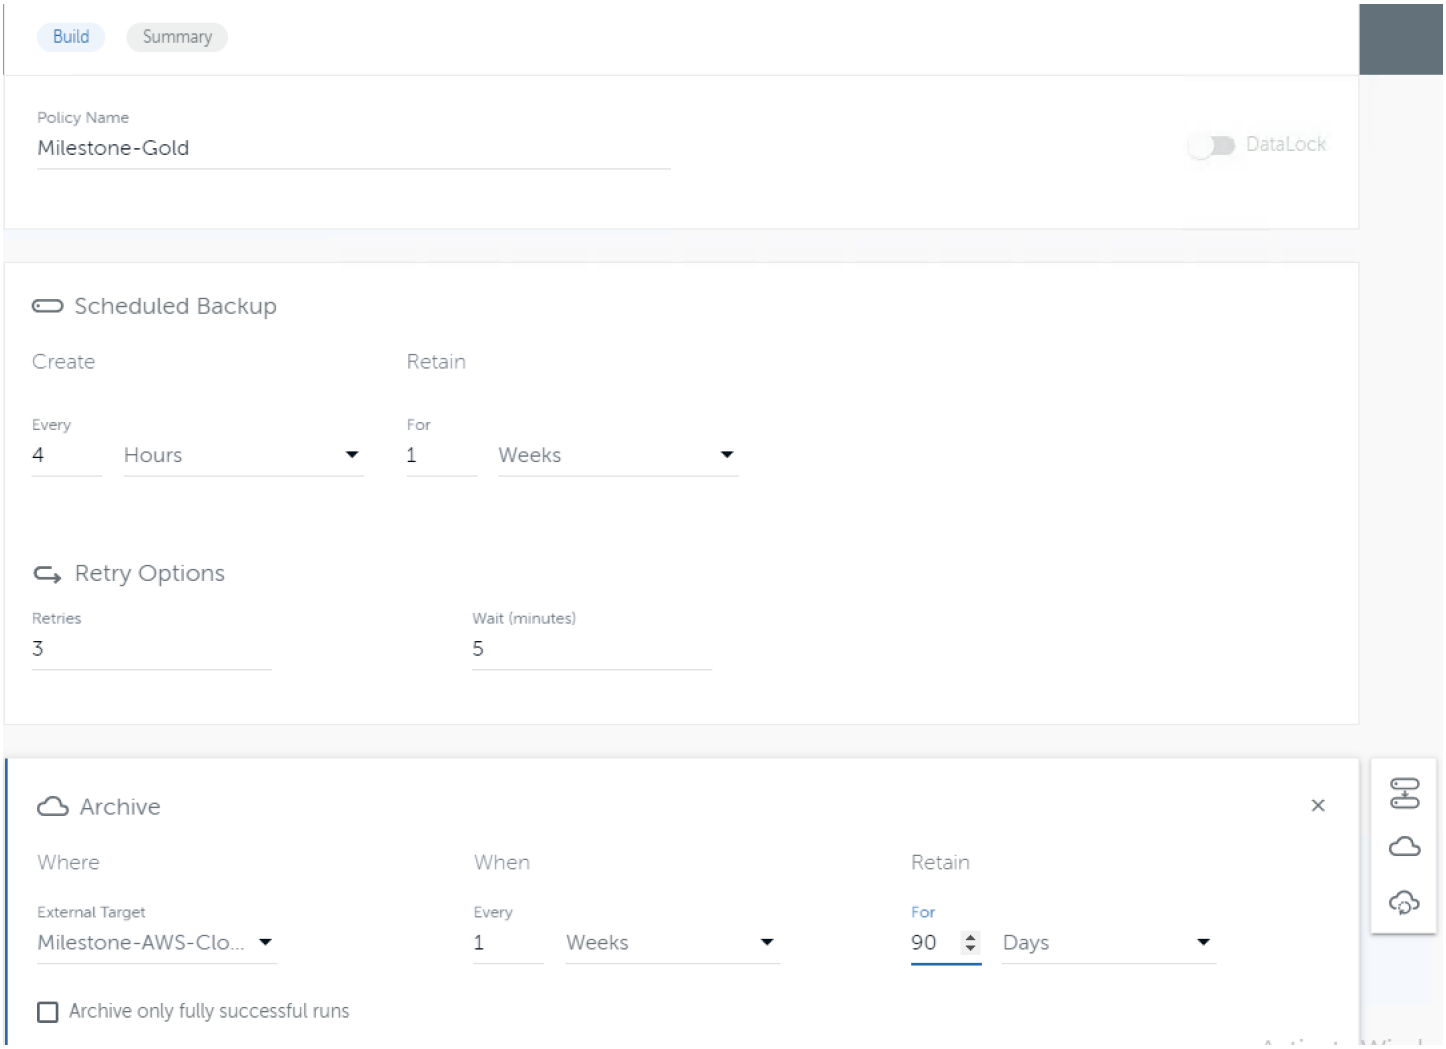

Figure 10.
Configuring protection policy

Protection jobs define one or more groupings of virtual machines or views for protection that comply with a specific protection policy. Jobs also set the time and time zone of the job and the storage domain that will contain the protected data (Figure 11).

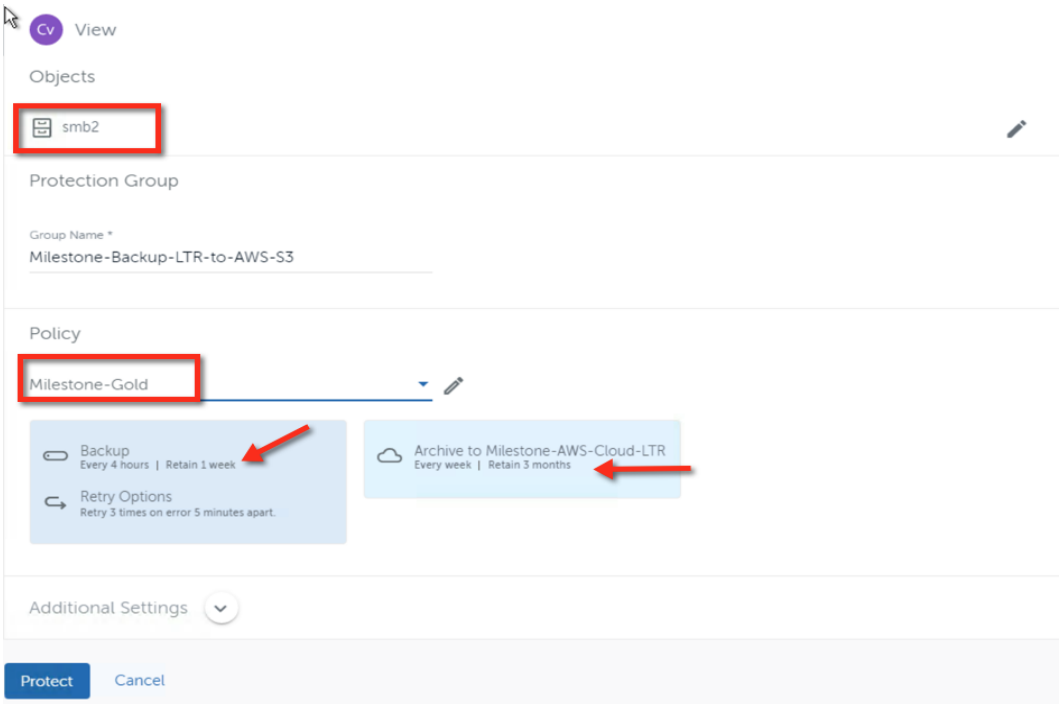

Figure 11.
Configuring a protection job

Policies are usually created to ensure RPO compliance and are made available to job protection across the entire cluster. Defined policies based around data protection requirements can be shared across protection jobs. Implement a naming scheme that allows easy identification of policy attributes.

As seen above, we also leveraged Cohesity to simplify data management and configured it to provide snapshot-based data protection for the SMB share for Milestone Archive DB datasets. Not only does this snapshot-based data protection provide multiple advantages from a recovery perspective but also allows for data sets to be cloned to provide multiple users a historical view into the Milestone datasets.

As shown in Figure 11, the SMB share was configured for data protection with the Milestone-Gold policy to use Cohesity snapshots, and the Milestone data set was retained locally on the cluster for 1 week. Then as part of the same policy archiving was configured for these data sets, copying them to AWS S3 blob storage for long-term retention for 3 months.

Cohesity simplifies data management by enabling long term retention or archiving of the Milestone data set by providing native capabilities to archive data to the cloud.

The AWS S3 external target was configured with the settings shown in Figure 12, which support overall data efficiency and provide data security with AES-256 encryption, compression, source-side deduplication, and if required, advanced features such as bandwidth throttling.

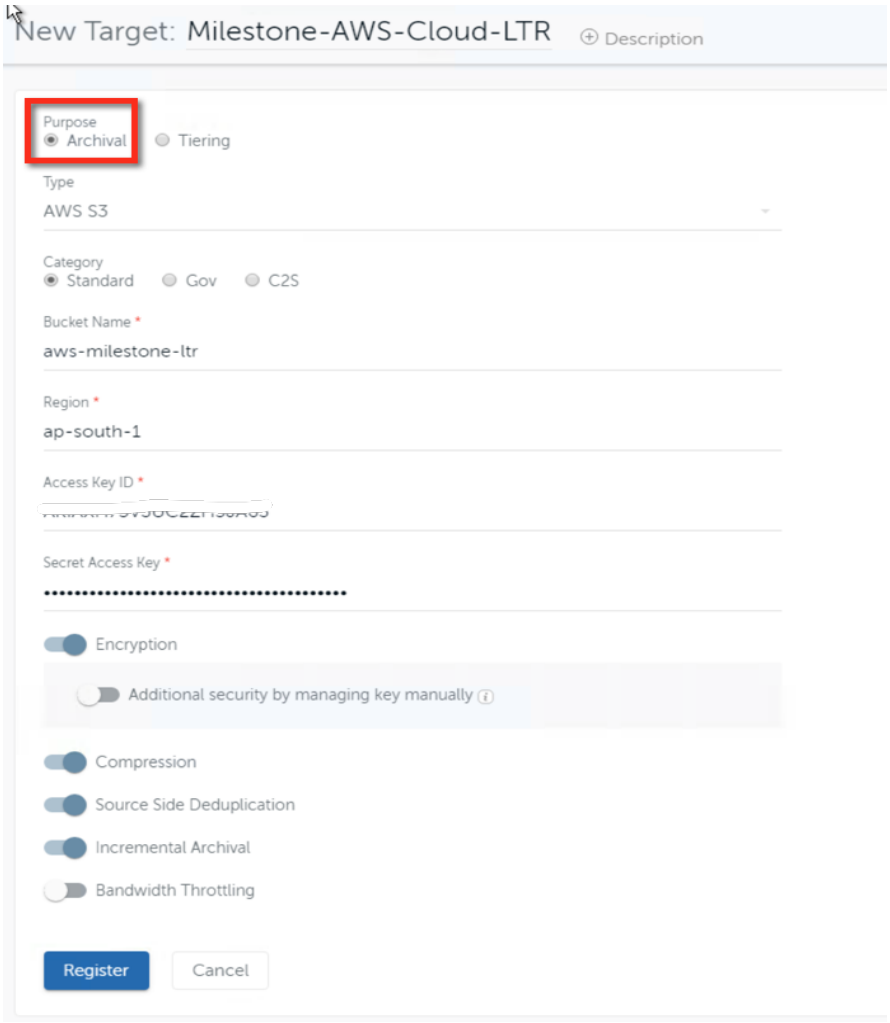

Figure 12.
Configuring an external storage target

Scaling the solution

This solution is easily scalable from a software and hardware standpoint as the solution incorporates the best of hyperconverged architectures both in the primary with HyperFlex and data management platform with Cohesity. When more cameras are added to the network, storage nodes can easily be added to the existing solution. To improve the decoding capability and performance of the system running XProtect Smart Client and for data analytics, you can add a computing node with a graphics processing unit (GPU).

Performance analysis

The video feed to the recording server was generated using the feed server, which is connected to the top-of-the-rack (ToR) switch. The feed server is used by the StableFPS driver, which is installed on recording server to generate network load. System performance was measured by using high-resolution 1920 x 1080 H.264 codec video that contained a section with motion in the first 20 to 25 percent of the total video followed by video with no motion for the rest of the video. The number of frames was set to 30 frames per second (fps).

The test was run for 7 days, and the results were captured. During the test, on each recording server, ingress data throughput (data coming into recording server media) was 50 MBps, but the average disk throughput of each recording server virtual machine was 12.2 MBps, because of the use of the Cisco HyperFlex inline deduplication feature. The total throughput was 183.0 MBps, and the average ESXi CPU use was 77.7 percent for 1500 camera simulation tests. Also, during the 7-day test, latency, vCPU use, memory use, and network throughput were all well within the limits.

Figure 13 shows the average read-write bandwidth and latency observed in Cisco HyperFlex Connect during the 1500-camera simulation test without archiving. The aggregate bandwidth of all 15 recording servers was 183 MBps with an average latency of 10.1 milliseconds (ms), and each virtual machine contributed an average bandwidth of 12.2 MBps.

Figure 13.
Average read-write throughput and latency of all recording servers observed in Cisco HyperFlex Connect

Figure 14 shows the latency and bandwidth of Cisco HyperFlex storage during the archiving of a single recording server. During the archive process, the average read bandwidth increased without affecting the write throughput. The average read bandwidth was 292 MBps with an average latency of 6.3 ms; the average write latency was 16.2 ms.

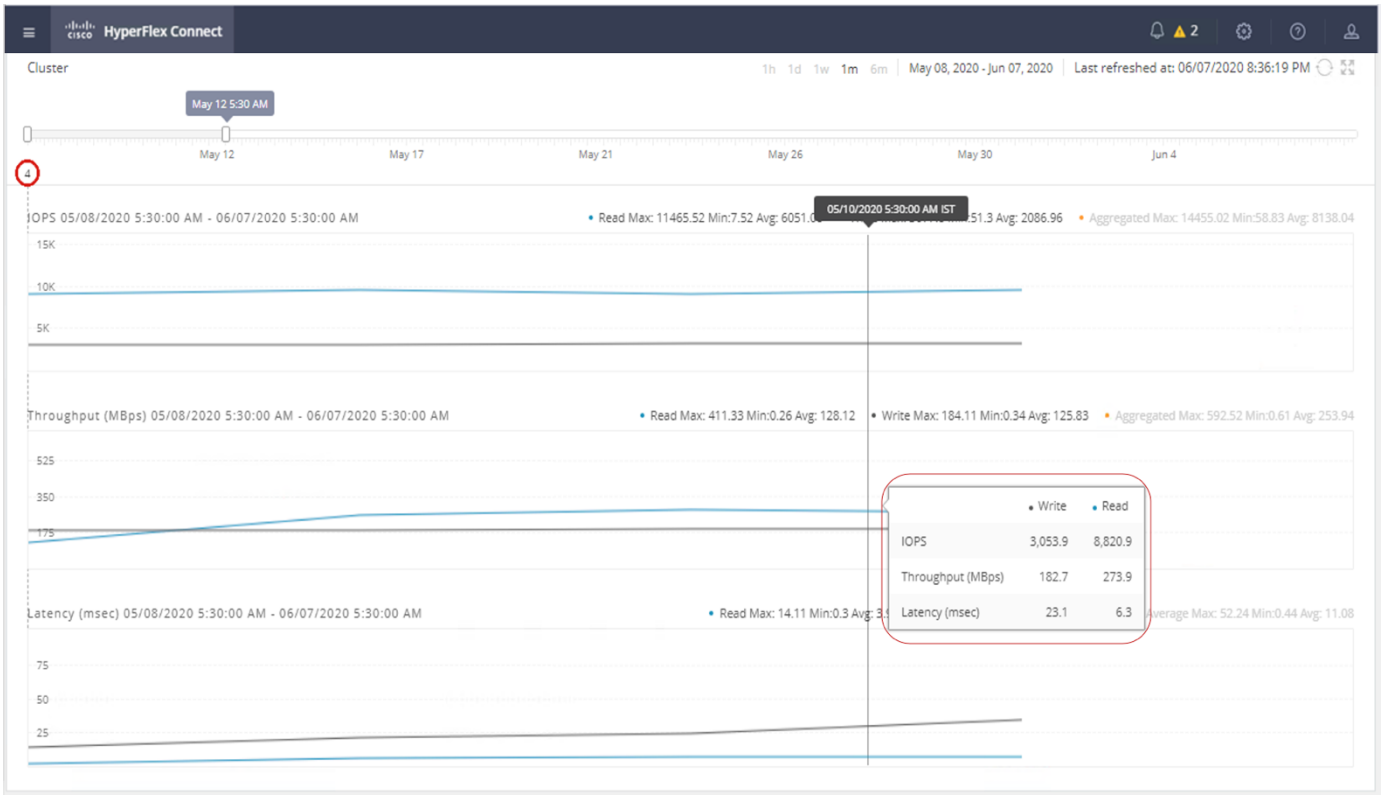

Figure 14.
Bandwidth and latency of Cisco HyperFlex storage

Figure 15 shows the latency and bandwidth on Cohesity DataPlatform observed during the archive process.

A single recording server with 100 cameras generated about 1 TB of data for 24 hours. This data was archived to a single node in the three-node Cohesity cluster at a rate of close to 290-300 MBps. Because the system is a scale-out system, Cohesity performance scales linearly based on the load being distributed across all the nodes in the cluster.

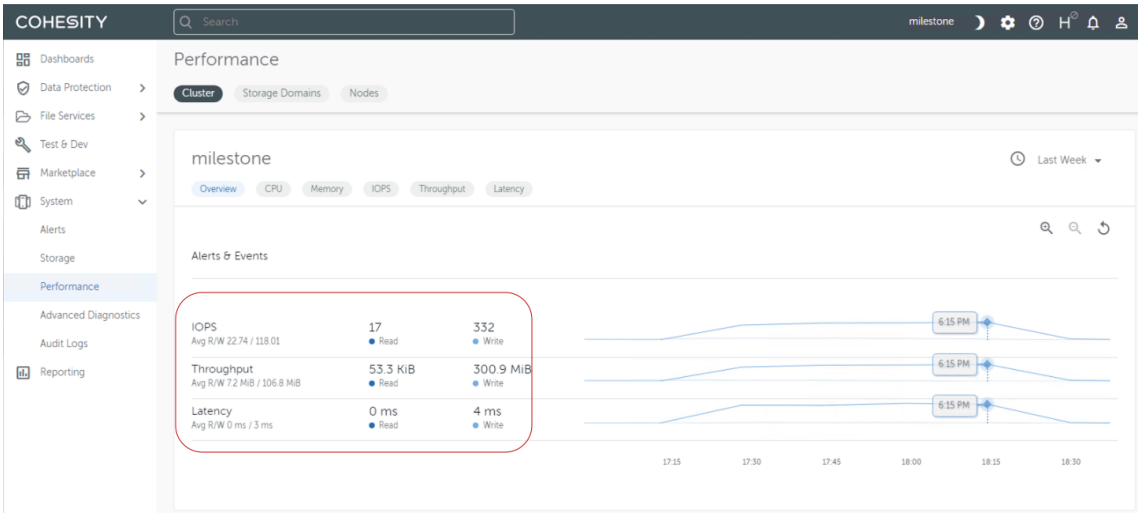

Figure 15.
Bandwidth and latency of Cohesity DataPlatform

Figure 16 shows the average vCPU use of all 15 recording servers, and Figure 17 shows the average ESXi host CPU use for the 7-day test. The average vCPU usage of all recording servers was 63.3 percent without any media loss. Also, during the 7-day test, latency, vCPU use, memory use, and network throughput were all well within the limits.

Because only 15 recording servers were configured, one of the ESXi CPUs was lightly loaded. This available resource can be used to deploy infrastructure virtual machines and Milestone VMS components.

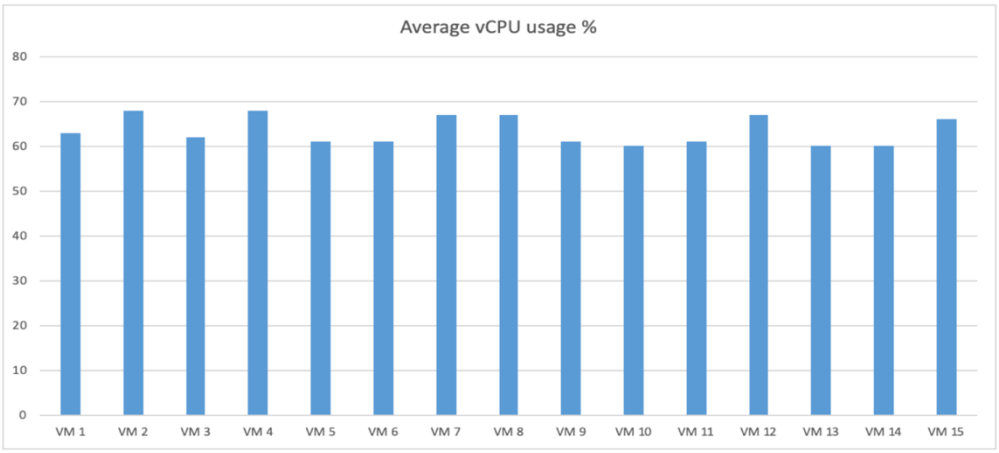

Figure 16.
Average vCPU use of all recording servers

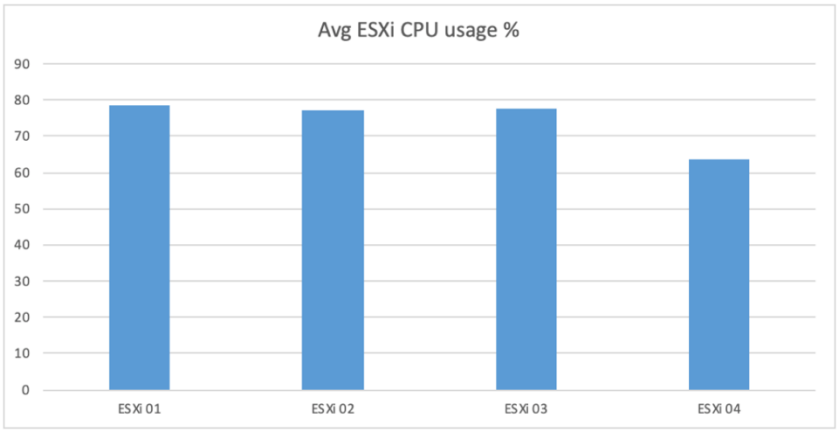

Figure 17.
Average ESXi CPU usage

Sizing

The XProtect recording server, which records all the camera video streams, is a resource-intensive application that continuously puts a high load on CPU, storage, network, and memory resources. Milestone recommends that, for a stable system, network, CPU, and memory use should not exceed 70 percent.

When CPU use averages more than 70 percent, the XProtect recording server can become overworked, leading to an increase in buffer overflows and dropped frames. However, this is not the only cause of dropped frames. Disk I/O can also cause frame loss, so XProtect has special counters on the OS that

monitor for dropped frames. Disk read latency is a concern because it causes visible latency when operators are watching live and recorded video. Visible latency is unacceptable for most customers.

Hence, designing CPU, memory, storage, and networking resources for performance-intensive workloads and applications in a virtualized environment is crucial. Otherwise, resource constraints will significantly affect application performance. This section provides sizing guidelines for live database recording.

These are the limits expected in a stable system:

- Network loads (receiving or sending) should not exceed 70 percent of available bandwidth.
- Processor load should not exceed 70 percent.
- The percentage of committed bytes in use should not exceed 70 percent.
- Disk latency (reading or writing) should not exceed 200 ms.

CPU and memory

The video resolution, fps rate, and video codec require more processing power, and the number of cores and CPU frequency affect the performance of the system. The CPUs used for Milestone certification were Intel Xeon Gold 6248 processors with 20 cores at 2.5 GHz. Each recording server was configured with 10 vCPUs and 32 GB of RAM. The average vCPU use was less than 65 percent, and the average memory use was less than 25 percent for the 100-camera simulation test.

The CPU and memory configuration can be changed based on the system sizing to meet the actual requirements. GPUs can be deployed on Cisco Hyperflex nodes to offload some of the encoding functionality from the recording server virtual machine. You can configure a dedicated compute-only node with a GPU to handle video analytics or deploy an XProtect Smart client to use the GPU to offload encoding workload. This configuration increases the number of streams that can be viewed on video walls.

Storage

Table 5 shows the storage space required by a single recording server and by all 15 recording servers. The calculation of disk space consumed is based on a video stream size of 4 Mbps per camera.

Table 5. Storage space calculation

Video Stream Bandwidth (Mbps)	Number of cameras per Recording Server	Ingress data (MBps)	Egress data to storage with inline Dedup enabled (MBps)	Storage required by Recording Server per day (TB)	Total Recording Servers	Total bandwidth (MBps)	Total storage required by 15 Recording Servers per day (TB)
4 (0.5 MBps)	100	50	12.2	1.01	15	183	15.1

The bandwidth requirements will change with the video and camera resolution and frame rate, and you should size storage accordingly. In calculating the storage space requirements, you need to consider the use of Cisco HyperFlex deduplication and the replication factor. RF3 will keep three redundant replicas of the data. This configuration consumes more storage resources, but it helps ensure better protection for your data in the event of a node or disk failure. A replication factor of 2, in contrast, will keep two redundant replicas of the data. This configuration consumes fewer storage resources, but it reduces the capability of the data protection function to tolerate only one node or disk failure. However, for higher storage requirements, you can configure the Cisco HyperFlex system with RF2.

Cohesity resiliency

Cohesity provides a fully distributed architecture. Every node in the cluster runs the same software. Cohesity software operates as a set of cooperating services, with each performing a specific function and providing a single application. These services communicate with each other within the same node, and

they also communicate with services on the other nodes in the same cluster, providing a unified set of features. An advantage of this architecture is that a defect in one service doesn't affect any other service, and improvements to each service can be made independently of the others as long as the communication interfaces remain the same. Each software component within Cohesity is built with a focus on resilience.

These are some of the features of the Cohesity architecture:

- Strong consistency: A read operation always returns the most recently written value. Consistency is achieved using the Paxos algorithm. Reading the same object on different nodes by different clients is guaranteed to return the same value. Strong consistency is maintained even during node upgrades.
- Consistent hashing: Data is spread across the cluster in a uniform way using a consistent hashing algorithm.
- Full distribution: The design has no single point of failure. Data is distributed on other nodes according to the specified replication factor and erasure coding or policies.
- Self-balancing properties: Data is redistributed automatically when nodes are added to or removed from the cluster. Redistribution is fast, because very little data movement is required for rebalancing.
- No single point of failure: When a node or disk fails, another node or disk will serve read and write processes. When the failed component is restored to use, it is automatically healed through seamless integration.
- No disruption on failure: When a failure occurs, new write operations continue to be processed with the same redundancy levels. Session state information is retained so that running jobs are not interrupted, even during node software upgrades.
- No disruption on upgrade: Cohesity's clustered architecture allows sequential node-by-node upgrades. Nodes can be upgraded to the latest version without any interruption of client I/O processing. Any active I/O access through the virtual IP address is transferred as the virtual IP address is seamlessly transferred to another node.

Cohesity provides configurable resiliency on HDD and node failures. Both RF2 and RF3 along with the erasure coding scheme are supported with the Cohesity SpanFS file system.

The replication factor specifies the number of replicas of a unit of data. The unit of replication is a chunk file, and a chunk file is mirrored to either one or two other nodes, depending on the replication factor value used. RF2 provides resilience against a single data unit failure, and RF3 provides resilience against failure of two data units.

Erasure coding is a scheme though which a number of usable data stripe units are protected from failures using code stripe units, which are in turn derived from the usable data stripe units. A single code stripe unit can protect against one data (or code) stripe failure, and two code stripe units can protect against two data (or code) stripe unit failures.

Depending on the resiliency and fault tolerance values used, the ratio of raw-to-usable capacity will vary. Figure 18 provides an overview of the usable capacity when different replication factor and erasure coding schemes are used.

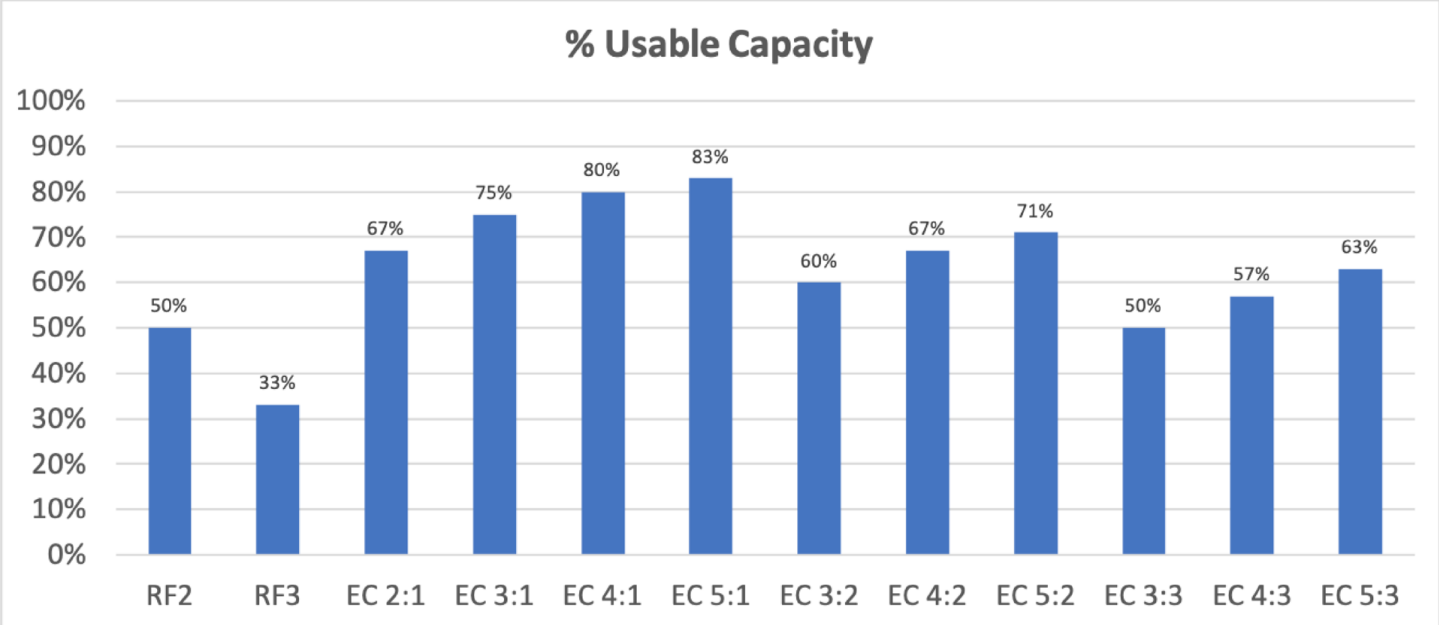

Figure 18.
Usable capacity based on replication factor and erasure coding

For additional information, see the Cohesity resilience white paper at https://docs.cohesity.com/HomePage/PDFs/Cohesity-White-Paper-Fault-Tolerance-Data-Integrity.pdf.

Network

Network topology and bandwidth have a significant impact on system performance and stability. Network traffic should be separated using VLANs and a dedicated network adapter.

The network throughput for the 1500-camera test was 765 MBps. Network bandwidth and topology should be designed properly for a larger system environment.

Conclusion

Cisco, Milestone, and Cohesity provide a compelling solution that brings together the best of next-generation technologies for video surveillance use cases for large enterprises and smart city solutions. Governments around the world are using the power of video surveillance to make the lives of their citizens safer and more productive and to enhance the quality, performance, and interactivity of urban services.

The validated solution described in this document is essentially infinitely scalable in its software and hardware because it incorporates the best of hyperconverged platforms. It uses the Cisco HyperFlex platform for the primary system and Cohesity DataPlatform for backup and unstructured data on one unified architecture based on Cisco UCS. It simplifies data management from on-premises to multicloud hybrid environments and scales seamlessly to accommodate data growth. All this is achieved with lower total cost of ownership (TCO), simplified operations, and a pay-as-you-grow consumption model.

This solution allows enterprises to easily manage and extract insights from data. It prepares them on their journey to leverage artificial intelligence (AI) and machine learning and other advanced technologies to gain more value out of their data.

References

For additional information, please see:

- https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html
- https://www.cisco.com/c/en_in/products/hyperconverged-infrastructure/index.html
- https://docs.cohesity.com/6_5/Web/UserGuide/Content/Welcome/Welcome.htm
- https://doc.milestonesys.com/sysarch/pdf/2020r1/en-US/MilestoneXProtectVMSproducts_SystemArchitectureDocument_en-US.pdf

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Printed in USA 220134.5 07/20