

Virtualized Computer Vision for Smart Transportation

Based on Dell Infrastructure with Milestone Systems

H19293

Abstract

This design guide describes the Dell Validated Design for Computer Vision applications with Milestone Systems XProtect. The design is presented and validated with Milestone Systems, BriefCam, and Ipsotek applications.

Dell Technologies Solutions



Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Chapter 1: Introduction.....	5
Solution introduction.....	5
Document purpose.....	5
Chapter 2: Architecture concepts and requirements.....	6
Architecture concepts.....	6
Solution requirements.....	6
Chapter 3: Solution Architecture.....	7
Physical architecture.....	7
Virtualized infrastructure.....	9
Computer vision applications.....	9
Milestone architecture.....	10
BriefCam architecture.....	15
Ipsotek architecture.....	19
Chapter 4: Validation.....	23
Validation overview.....	23
Performance validation.....	24
VM placement with DRS.....	24
Placement constraints.....	24
DRS Rule Details.....	24
VM placement for validation.....	25
Video workload.....	26
Camera mapping to Milestone.....	26
Performance results and findings.....	27
Milestone XProtect performance tests.....	27
BriefCam performance tests.....	29
Ipsotek performance tests.....	31
Full system performance tests.....	33
Performance summary.....	34
High availability validation.....	34
VxRail high availability.....	35
High availability results and findings.....	36
BriefCam HA results.....	36
Ipsotek HA results.....	37
Milestone HA results.....	38
High availability summary.....	41
Chapter 5: Sizing the solution.....	42
Overview.....	42
Sizing guidelines.....	42
System sizing.....	43
Storage sizing.....	43

Cameras per node.....	43
Scaling guidelines	43
Chapter 6: Backup and restore operations.....	45
BriefCam.....	45
Ipsotek.....	45
Milestone Systems.....	46
Chapter 7: Monitoring the solution.....	48
VxRail platform monitoring.....	48
BriefCam.....	48
Ipsotek.....	48
Milestone Systems.....	49
Chapter 8: Troubleshooting recommendations.....	50
BriefCam.....	50
Ipsotek.....	50
Milestone Systems.....	51
Chapter 9: Summary and conclusions.....	52
Appendix A: References.....	53

Introduction

Topics:

- [Solution introduction](#)
- [Document purpose](#)

Solution introduction

Transportation organizations across the globe are turning to computer vision capabilities powered by artificial intelligence (AI) to create a safer and more engaging passenger experience with improved security, and advanced safety measures such as curb-to-gate touchless check-in.

Dell Technologies' computer vision solutions enable transportation organizations to provide a more seamless passenger experience by bringing together the right combination of AI, computer vision technologies, and operational workflows. Our full orchestration of intelligent edge video, compute, storage, networking, analytics, and cloud integration technologies deliver an end-to-end capability that enables organizations to focus on what matters most - the passenger. In addition to streamlining the passenger experience, our solutions allow organizations to realize greater business value, improved safety, and security, plus increased operational savings and efficiencies.

Document purpose

There are multiple options to choose from when designing an integrated computer vision (CV) and video management system (VMS) for a large environment. Most implementations include products from multiple vendors and then require a significant integration project effort before operations can begin. Dell Technologies together with three market-leading VMS and CV vendors have designed an integrated solution running on a Dell Validated Design for AI platform that is easy to operate and easy to grow. This document describes how we designed and tested our integrated solution on a common platform and presents some preliminary sizing guidelines that can jump-start an environment-specific design tailored to the needs of a large site.

This Design Guide describes how to implement virtualized instances of Milestone XProtect for video management together with CV applications from both BriefCam and Ipsotek. All three applications were implemented following guidance from the vendors on how to take advantage of native high availability (HA) features and scale-out performance that would be critical in an enterprise-class solution.

We chose the VxRail HCI system with VMware virtualization plus Nvidia A40 server-class GPUs. NVIDIA vGPU compatibility with the Dell Validated Design for AI platform provides better resource utilization and flexibility when implementing large systems with many virtual machines. The VxRail systems and VMware virtualization are design components that are already familiar to many IT and OT (operational technology) professionals. These components also provide additional features for managing high availability and scale-out needs that can complement the native features of the CV and VMS applications.

All VMS and Computer Vision products were installed on a single VxRail 5-node cluster and tested to demonstrate the value of a common platform. Storage for the video management archive functionality was provided by Dell PowerScale scale-out NAS storage. All video streams used for scale-out testing were provided by camera simulation software reading from prerecorded local video files in a continuous loop. The VMS high availability testing was performed with physical cameras since our simulation software did not support dynamic rerouting of output streams. We tested simulated camera streams being archived to the VMS simultaneously with two CV applications performing real-time analytics on a subset of the simulated camera streams.

This document confirms that the VxRail with VMware virtualization and NVIDIA GPUs provides a reliable and scalable Dell Validated Design for AI platform for hosting VMS and CV applications from different vendors on a common platform.

Architecture concepts and requirements

Topics:

- [Architecture concepts](#)
- [Solution requirements](#)

Architecture concepts

The Dell Validated Design for AI is a multipurpose platform focused on the display and analysis of CCTV video. The architecture uses platform components from Dell, VMware, and Nvidia that are widely used by IT organizations around the world. This solution adds additional value to organizations that want to host a curated and validated set of video management processing and GPU accelerated analytics engines on a common platform.

This design guide covers a selected set of important computer vision workloads that are essential to the safe and efficient management of airports. This release is intended to demonstrate the feasibility and value of using a common platform for all transportation computer vision workloads that both serves the needs of operational IT and is also well known to organizational IT decision-makers.

This Dell Validated Design is based on a VxRail GPU optimized HCI platform leveraging virtually provisioned GPU computer resources that are used for video inferencing and visualization. Our purpose in offering this solution is to help organizations shift from an appliance-per-application-based approach to a common virtualized approach with the simplicity of a single hyperconverged compute and storage management for many applications.

VMware vSphere Distributed Resource Scheduler (DRS) is a feature that can improve application service levels by guaranteeing appropriate resources to virtual machines including deploying new capacity to a cluster without service disruption. DRS allows system administrators to take advantage of the intelligent placement of VMs for all three applications across the VxRail cluster that enhanced the overall system reliability and cluster resource usage. Staff responsible for Day 2 operations also benefit during planned and unplanned operations through the use of affinity and anti-affinity rules with DRS to maintain the most resilient configurations and maximum uptime.

Solution requirements

Dell Validated Design

- Demonstrate the integration of Milestone XProtect VMS, BriefCam CV, and Ipsotek CV hosted on a common platform.
- Validate that multiple Nvidia GPU accelerated applications can be hosted on a common platform.
- Recommend optimal use of storage technologies and configurations for each application.
- Provide backup and restore recommendations for all applications.
- Provide monitoring and alerting best practices.
- Document any troubleshooting recommendations discovered during development and testing.
- Develop sizing guidelines that support quote development for small, medium, and large configurations.

System reliability

- Design and test a highly available architecture that customers would use for mission-critical CV tasks.
- Conduct failover tests for all CV applications that leverage vGPUs.
- Validate that VMS failover in virtualized environments is consistent with previous bare-metal testing

Solution Architecture

Topics:

- [Physical architecture](#)
- [Virtualized infrastructure](#)
- [Computer vision applications](#)

Physical architecture

Compute

In order to meet the requirements for performance, maintainability, and horizontal scalability, the solution is built using a VxRail V670F 5-node cluster. This V Series VxRail is a 2U platform with Nvidia GPU support. It is also configured with all flash drives for optimal performance. The full specification is available [here](#). The design testing was performed with three nodes that were dedicated to video processing while running three market-leading applications simultaneously. Milestone XProtect Systems provided the VMS functionality for our solution and Ipsotek and BriefCam provided the CV functionality products. In a full production environment, there are going to be trade-offs that need to be considered when designing the initial placement and any use of automation for maintaining reliable operations. This guide describes what we validated in our lab environment. Please consult with all software suppliers and any system integrators prior to implementing a production-grade safety and security system.

The base Hardware specs are as follows:

Table 1. Base hardware specs per node

CPU	2 x Intel(R) Xeon(R) Gold 6354 CPU @ 3.00 GHz (72 vCPU)
Memory	512 GB
Storage	2 x 800 GB SSD cache drives 8 x 3.84 TB SSD capacity drives
GPU	2 x A40 48 GB Nvidia GPU

Network

The following table describes the networks that are configured as part of the validated design:

Table 2. Networks for the validated design

Network	Description
vSphere Management	Used by ESXi for host management.
vMotion	Used by ESXi for vMotion.
vSAN	Used by ESXi for vSAN traffic.
User	Used for user access to applications running on the VxRail.
Camera	The camera network isolates camera traffic so only specific applications receive video streams.

Table 2. Networks for the validated design (continued)

Network	Description
Storage	To optimize performance, application access to data storage is gained using a storage-only network.

By configuring separate networks for the different traffic types, we get an additional layer of security and performance improvements.

Storage

The solution design uses two storage technologies:

- VMware vSAN for virtual machine disks, and Milestone Tier 1 storage
- Dell PowerScale for VMS video archiving

vSAN for VM local disks

The Dell VxRail cluster was configured with vSAN datastores that are available for local disk mapping by all the VMs running on the cluster. The vSAN storage was configured with a RAID 5 storage policy to provide fault tolerance for all the applications writing to vSAN storage. This protection level maximizes the storage capacity while providing protection against a node failure.

The vSAN storage was used for all disks needed by Milestone Systems XProtect, Ipsotek, and BriefCam VMs including the operating system and any local file locations. This includes all processing and management VMs needed by each of the software solutions.

The vSAN was also used for Milestone Systems Tier 1 storage. This storage is written to vSAN storage and then archived to PowerScale long term storage.

PowerScale for bulk video stream archive

Dell PowerScale is used to store large volumes of video data that must be safely retained based on the needs of the customer. An A3000 PowerScale appliance was used in this design for the video stream archive. Shares on the PowerScale A3000 were set up as Continuous Availability (CA) shares. The SMB network protocol is used between the Milestone Systems XProtect and the PowerScale A3000. File shares are defined as CA shares and use the SMB3 Witness protocol. The witness protocol enhances client failover for SMB3 CA shares. Witness notifies the Expansion Server when a PowerScale node becomes unavailable without the need to wait for the SMB3 connection time out. Other PowerScale models including the A2000 can be used with this solution. This storage is also responsible for responding to requests for retrieving historic video data for analysis by CV applications.

For more information about Dell storage options with Milestone, see [Configuration Best Practices-Dell EMC storage solution with Milestone XProtect Corporate](#).

Architecture

The high-level architecture is as follows:

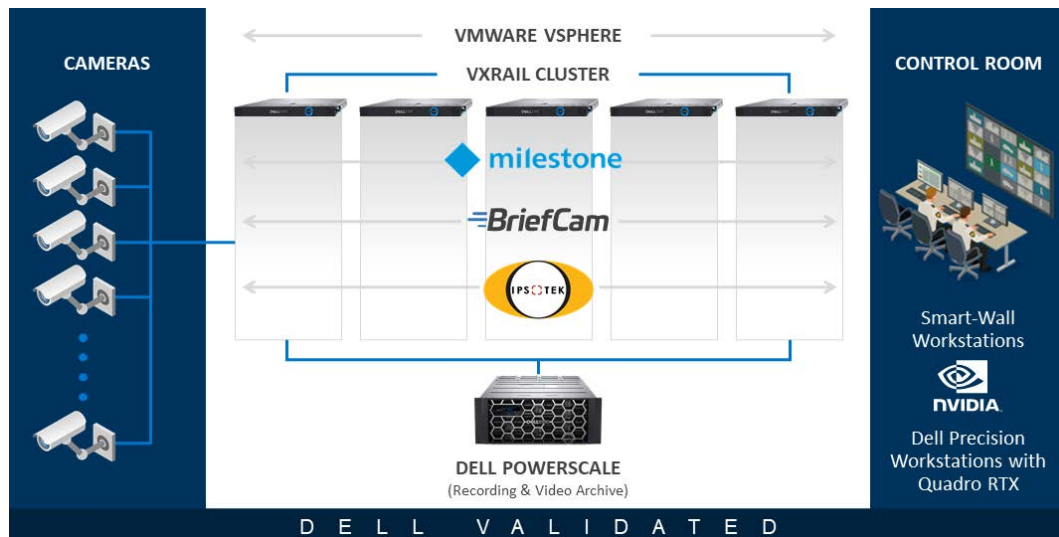


Figure 1. Dell Validated Design high-level architecture

This design augments the HA support provided by each application by placing selected VMs and services for each vendor across the cluster. Some components of a complete VMS and CV solution, such as the cameras and Control Room shown the previous image are not described in this design document, but can easily be integrated with the solution.

Virtualized infrastructure

This solution is based on the [Dell Validated Design for AI](#) which is a platform that is jointly engineered by Dell, VMware, and Nvidia. The Dell Validated Design for AI is an enterprise grade platform that is fully validated by Dell, VMware, and Nvidia. The design supports virtualization with GPUs and runs on Dell PowerEdge and VxRail HCI infrastructure. VxRail was selected for Computer Vision due to its ability to scale to support thousands of cameras.

The main components from the Dell Validated Design for AI that are leveraged are:

- VMs - All VMS and CV systems were fully virtualized.
- vGPU - Virtualized GPU capability that is required to process analytic camera streams.
- vSAN - All VMs are protected with vSAN storage.
- vCenter - All system management and operations are performed through VMware vCenter.

Computer vision applications

The primary goal of developing and testing this solution design was to determine if a single-converged platform could be used to virtualize VMS and CV applications from multiple vendors. The initial applications that were used for validation include Milestone Systems XProtect VMS, Ipsotek, and BriefCam.

Milestone Systems XProtect VMS

The VMS system software and versions tested were:

- Windows Server 2019 Operating System
- XProtect Corporate 2022 R1
 - Management Servers
 - Recorder Nodes
 - SQL Server Databases

BriefCam CV

The BriefCam solution software and versions tested were:

- Windows Server 2019 Operating System
- BriefCam 6.3
 - PostgreSQL Database
 - BriefCam Processing Servers (with GPU)
 - BriefCam Web Services

Ipsotek CV

The Ipsotek solution software and versions tested were:

- Windows Server 2019 Operating System
- Ipsotek 11.7.1
 - Database Nodes
 - Ipsotek Management Nodes
 - Ipsotek Processing Nodes (with GPU)

Milestone architecture

Milestone Systems develops products for large-scale facilities with high-security requirements. The XProtect family of products ensures end-to-end protection of video integrity while maximizing hardware performance.

The XProtect VMS products provide video management capability for environments covering a wide range of use cases and scales. Milestone Systems offer four versions of XProtect:

- XProtect Corporate
- XProtect Expert
- XProtect Professional+
- XProtect Express+

This suite of versions can support applications that range from protecting individual stores from vandalism to managing a multi-site, high-security facility. All solutions offer centralized management for all devices, servers, and users, including a flexible rules processing system driven by schedules and events. Our development work for this Design Guide was performed using the XProtect Corporate version.

The remainder of this document is specific to the XProtect VMS 2022 R1 Corporate version with significant changes to the system architecture and many new features not available in earlier versions. Implementations of the XProtect Corporate VMS system will typically consist of the following main components:

- The central management servers
- XProtect Download Manager
- One or more recording servers
- One or more installations of XProtect Management Client
- One or more installations of XProtect Smart Client
- One or more users of XProtect Web Client and XProtect Mobile client (optional)

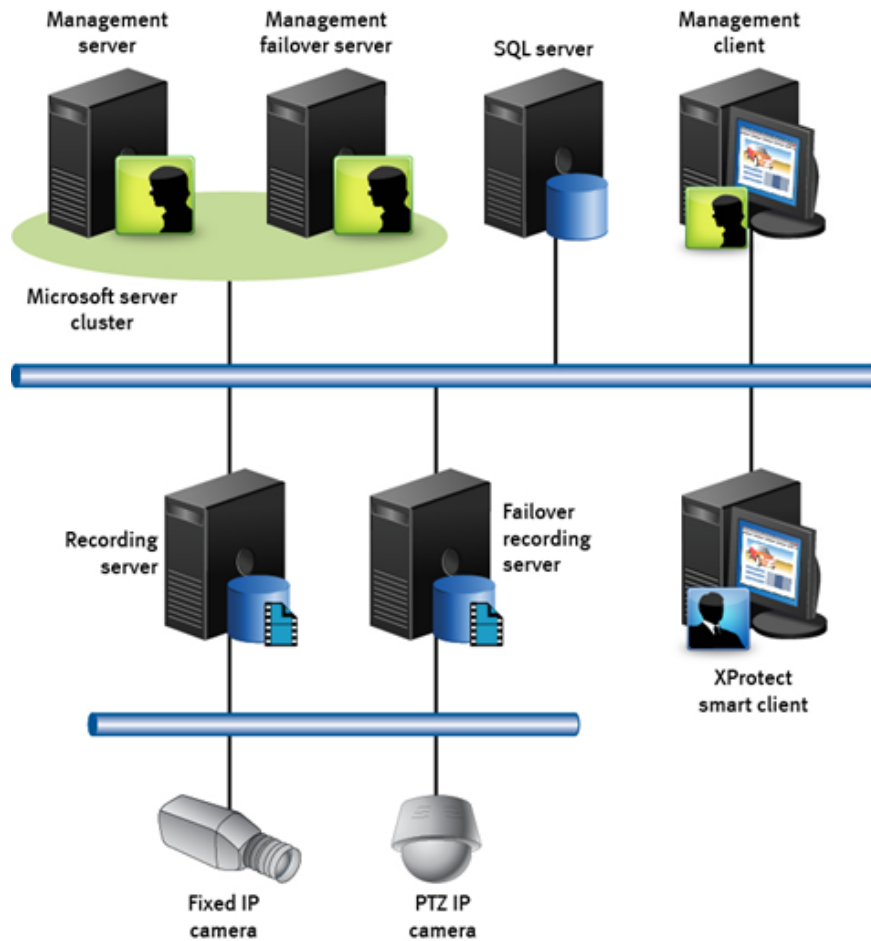


Figure 2. Milestone XProtect Corporate architecture

Scale-out architecture

To enable scaling up to thousands of cameras across multiple sites, the XProtect system consists of several components that handle specific tasks. For systems with more than 100 cameras, Milestone recommends that you use dedicated servers for all or some of the components.

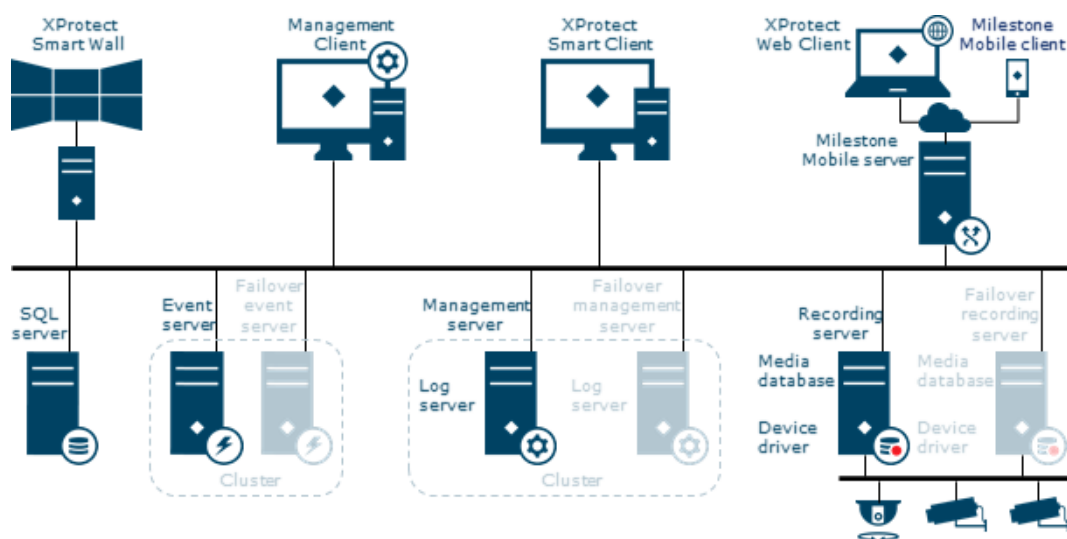


Figure 3. Milestone scale-out architecture

The full set of components shown above need not be available in all installations initially. Components such as failover recording servers or mobile servers can be added if the functionality they offer is needed at a later time for hosting and providing access to both XProtect Web Client and XProtect Mobile.

Depending on hardware and configuration, smaller systems with 50 to 100 cameras can run on a single server.

Configuration best practices

The following guidelines are based on previous internal testing and reports from customer field experience. None of these guidelines were exceeded during testing.

CPU	Do not assign more vCPUs to your VMs than the number of physical cores on the host machine.
Memory	The total amount of memory allocated to the VMs and the hypervisor should not exceed the total amount of physical memory available from the host.
Storage	<ul style="list-style-type: none">• Install Microsoft Windows and Microsoft SQL databases on vSAN devices according to the space and performance needs suggested by each application vendor.• For video recorder virtual machines, do not store archived video on the OS drive.• Do not use the OS drive for archived video.• Make the OS partition at least 120 GB.• For Tier1 storage, make use of dedicated drive from vSAN storage with Raid5 Thick Storage policy.• For Tier 2 archived video, configure PowerScale NAS share (e.g A3000) for video storage.
Networking	<p>When provisioning multiple archiving VMs on a host, do not exceed data transmission rates of 96 MB/s per node:</p> <ul style="list-style-type: none">• In-Guest iSCSI networks with other configurations might result in degraded performance.• Isolate video traffic on a different VLAN from vSAN storage traffic.• At least 10 GbE network cards are required (40 GbE preferred) for shared traffic links (management, video, and storage) configured with a Virtual Switch.• Alternatively, a dedicate1 GbE network card per VM for video traffic can be sufficient in some cases.• Network configurations that may result in multicast traffic being sent to all hosted VMs simultaneously should be avoided to limit the potential for negative performance impacts.

High availability

High Availability is an important consideration for security systems deployment to avoid loss of camera stream data during an unplanned outage. While no system can achieve 100% availability, the following components should be deployed with high availability considerations.

The following VMS components can run on a single virtual machine for small environments.

- Management server
- Event server
- Log server
- SQL Server

However, as noted in the Overview section, Milestone recommends hosting SQL Server on a dedicated server for applications with many devices and/or many event transactions. For dedicated SQL Server installations, the design must then consider fault tolerance for both the database and the three remaining services (management, events, and logging). We tested SQL Server installed with support for Always On Availability Groups (AAG) following the best practices documented by Microsoft. We dedicated two virtual machines for hosting two stand-alone (non-FCI) instances of SQL Server 2019. We created a virtual network name and virtual IP in Active Directory to support our AAG configuration. We followed this [AAG configuration](#) guide from Microsoft without modifications.

Management Server Failover

The management server can be installed on multiple servers within a cluster of servers. This ensures that the system has very little downtime. If a server in the cluster fails, another server in the cluster automatically takes over the failed server's job running the management server. The automatic process of switching over the management server services to run on another

server in the cluster takes up to 30 seconds. For more information about high availability options for the Management Server, see the Milestone Systems XProtect VMS documentation for [Multiple management servers \(clustering\) \(explained\)](#).

Recording Server Failover

A failover recording server has two services installed:

Failover Server service	Handles the processes of taking over from the recording server. This service is always running, and constantly checks the state of relevant recording servers.
Failover Recording Server service	Enables the failover recording server to act as a recording server.

- The failover server service checks the state of relevant recording servers as defined by settings in the FailoverServerConfig.xml file. The typical polling interval is every 0.5 seconds. If a recording server does not reply within 2 seconds and other conditions such as a successful ping of the management server are met, then the recording server is considered unavailable and the failover recording server takes over.
- Any failover solution will not provide complete redundancy but will provide a reliable way of minimizing downtime. When a failed primary recording server becomes available again, it will merge any recordings for recorders stored on the failover recording server back to its own storage. A loss of recordings during this process is very unlikely.
- Recordings stored by the failover recording server are automatically merged into the primary recording server's databases. The time it takes to merge, depends on the number of recordings, network capacity, and more. During the merging process, you cannot browse recordings from the period during which the failover recording server took over.
- In the event of a recording server failure, users should hardly notice that a failover recording server has taken over. A short interruption in service occurs, usually only for a few seconds, when the failover recording server takes over.
- During the service transfer exchange, users cannot access the video from the failed recording server. Users will be able to resume viewing live video as soon as the failover recording server has taken over.
- The recordings stored on the failover recording server can be played back for the period after the failover recording server took over. Clients cannot playback older recordings stored only on the failed recording server until that recording server is functional again and has taken over from the failover recording server.

Cold standby failover recording servers

In a cold standby failover recording server setup, you group multiple failover recording servers in a failover group. A cold standby failover group must contain at least one failover recording server.

The entire failover group is dedicated to taking over from any of several preselected recording servers if one of these becomes unavailable. You can create as many failover groups as needed to achieve the required level of outage protection. In a cold standby setup, the Failover Recording Server service is only started when the cold standby failover recording server takes over from the recording server. The time required to start the service depends on several factors including CPU characteristics, disk configuration, network performance, the number of camera devices configured, security settings and more.

When an active recording server fails, one of the available servers in the assigned primary failover group will take over. Grouping failover recording servers can provide a pool of servers in reserve that can protect one or more active recording servers. If the configured primary failover group contains more than one failover recording server at the time of failure, then one will be assigned to take over. You can specify a secondary failover server group that can supply servers to take over for a failed active recording server if all the recording servers in the primary group are in use. A failover recording server can only be a member of one failover group.

Recording servers in a failover group are ordered in a sequence. The sequence determines how the failover recording servers will take over from a failed active recording server. By default, the sequence reflects the order in which you have incorporated the failover recording servers in the failover group: the first in is the first in the sequence. You can change this order if needed.

If a failover recording server in a cold standby failover group must take over from another recording server during the merging process it postpones the merging process with the first failed recording server (Server A) and takes over for the newly failed recording server (Server B). When recording server B becomes available again, the failover recording server takes up the merging process and allows both recording server A and recording server B to merge back recordings simultaneously.

Hot standby failover recording servers

Hot standby failover recording server setup requires a dedicated failover recording server for each protected recording server. This one-to-one mapping allows the system to quickly transition a failover recording server from "standby" mode by synchronizing the correct or current configuration of the recording server it is dedicated to.

A hot standby failover recording server can take over much faster than a cold standby failover recording server. In a hot standby setup, the Failover Recording Server service is always running, allowing the hot standby server to take over faster than the cold standby failover recording server. Hot standby servers are assigned in a one-to-one pairing with recording servers. A failover group cannot also protect a primary recording server protected with a hot standby. You cannot assign failover servers that are already assigned to a failover group as hot standby recording servers.

A hot standby failover recording server only has to start its cameras on failover to deliver feeds since the Failover Recording Server service is always running. During the startup period, you cannot store recordings or view live video from affected cameras.

If a recording server with a hot standby server fails again during the merge back process, the hot standby server will take over again while retaining the recordings from the previous failed period. The hot standby recording server keeps all recordings until they are merged back to the primary recorder or until the failover recording server runs out of disk space.

A hot standby server cannot be configured with failover protection from either a standby failover group or another hot standby server.

VM designs

This section details the VMs used to host XProtect services for this testing:

Name	Cores	Memory	GPU	Storage	OS	T1 Storage	T2 Storage	Role
mil-db-1	6	12 GB	None	200 GB	Win Server 2019	None	None	Management SQL Server Database
mil-db-2	6	12 GB	None	200 GB	Win Server 2019	None	None	Management SQL Server Database 2
mil-mgt-1	6	12 GB	None	200 GB	Win Server 2019	None	None	Primary Management server
mil-mgt-2	6	12 GB	None	200 GB	Win Server 2019	None	None	Secondary Management Server
mil-rec-1 to mil-rec-12	6	12 GB	None	200 GB	Win Server 2019	2 TB Raid 5 vSAN Drive	20 TB , Power Scale Share	Recording servers (with PowerScale folder mapping)
mil-hot-1	6	12 GB	None	200 GB	Win Server 2019	2 TB Raid 5 vSAN Drive	None	Hot Backup Recording server
mil-cold-1 to mil-cold-2	6	12 GB	None	200 GB	Win Server 2019	2 TB Raid 5 vSAN Drive	None	Cold Backup Recording server

High level architecture

For our design, the Milestone XProtect VMs are distributed across the VxRail cluster by DRS based on affinity rules. The mapping process can be visualized as follows:

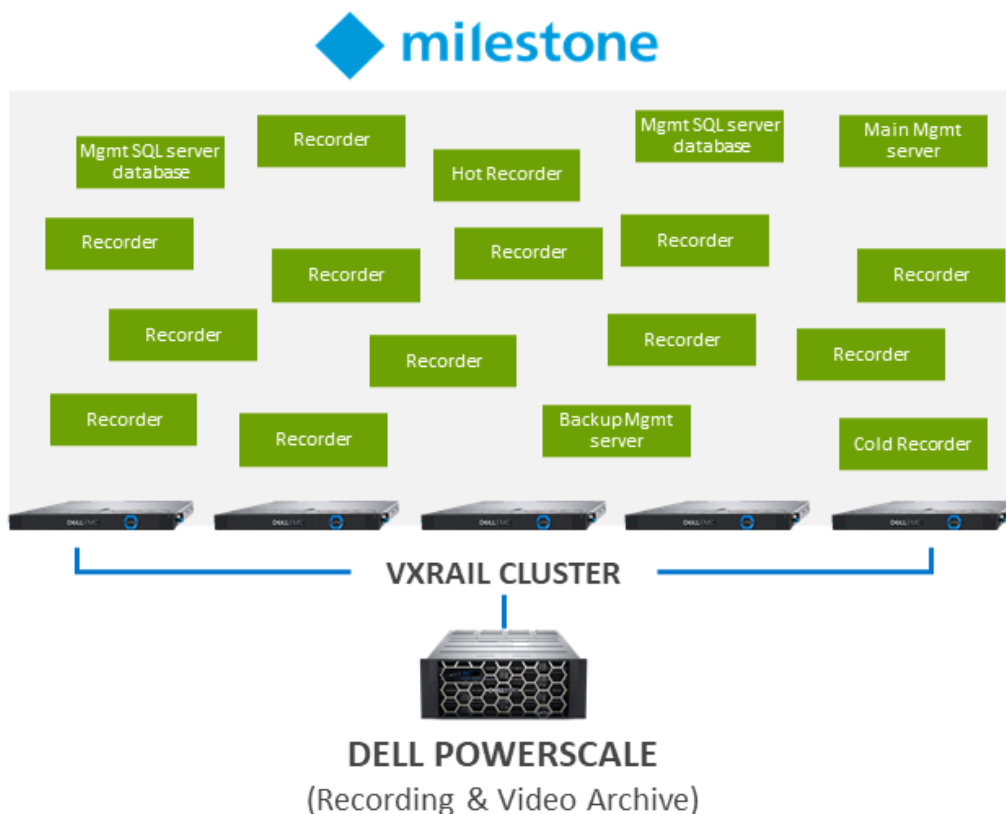


Figure 4. Milestone VMs across the VxRail cluster

BriefCam architecture

The BriefCam Computer Vision system is implemented using a microservices architecture that supports feature deployment and licensing flexibility. The features of the solution are packaged into a set of seven server roles that can be deployed in an all-in-one server configuration for proofs-of-concept or various multiserver configurations to meet performance and HA goals. BriefCam recommends deploying the application using an HA architecture to ensure business continuity even when some components fail.

A combination of well-known HA technologies is available for adding redundant nodes to a system that prevents a single point of failure (SPOF). However, adding redundant nodes does require extra hardware resources and increases deployment complexity. Designing the right solution for your production environment requires balancing cost, complexity, and the level of risk from a system outage.

Configuration best practices

BriefCam is driven by a microservices architecture, enabling scalability, flexibility, and redundancy. The full microservices implementation is described in the *BriefCam® Microservices Architecture* white paper. To obtain copies of BriefCam manuals, reach out to your BriefCam representative. The implementation developed for this design guide was a single-site environment. The following 14 microservices were enabled on servers configured for the design. However, not all were used during testing, such as LPR matching.

Table 3. BriefCam microservices

Microservice	Description
Alert Processing	The Alert Processing service is responsible for real-time video processing. The service must be hosted on a GPU-enabled processing machine.
Face Recognition	The Face Recognition service is responsible for the following activities: <ul style="list-style-type: none"> Monitoring the external watchlist folders for new face images

Table 3. BriefCam microservices (continued)

Microservice	Description
	<ul style="list-style-type: none"> Providing the aggregated status of the faces' feature-vector extraction process to the user, when uploading faces to a watchlist from the web UI
Fetching	The Fetching service is responsible for fetching the video footage from the VMS, storing it in the storage and making it available for BriefCam's Processing Server. By default, BriefCam's Fetching service uses two workers to fetch the videos from the VMS. The workers fetch five-minute video chunks.
Filtering	The Filtering service is responsible for handling in-memory object filtering for various scenarios in all the modules (REVIEW, RESPOND, and RESEARCH).
Face Recognition Matching	The Face Recognition Matching service is responsible for processing faces to find matches for filtering in the REVIEW module and for RESPOND alerts.
Maintenance	The Maintenance service is responsible for running BriefCam's automatic maintenance processes. For more information about maintenance, see the Maintenance and Data Retention section of the <i>BriefCam Administrator Manual</i> .
Lighthouse	<p>The Lighthouse service is the seed node of BriefCam's Akka cluster. Its main roles are:</p> <ul style="list-style-type: none"> Registering new services that join the cluster Providing service-discovery capability for all the other services in the cluster <p>BriefCam's Akka cluster consists of the following services:</p> <ul style="list-style-type: none"> LPR Matching BI Face Recognition Face Recognition Matching Filtering
LPR Matching	The License Plate Recognition (LPR) Matching service is responsible for processing license plates to find matches for filtering in the REVIEW module and for RESPOND alerts.
Notification	The Notification service is responsible for managing all aspects of notification and message delivery between the client application and server-based components.
Outbound API Gateway	This Hub service collects RESPOND alerts from the sites and if needed, sends them to a third-party service.
Rendering	<p>The Rendering service is responsible for:</p> <ul style="list-style-type: none"> Generating visual and video artifacts for the web client, such as rendering the synopsis videos and visual layers, exports and original videos. Validating uploaded video files before processing.
Video Streaming Gateway	The video streaming gateway is responsible for streaming the video artifacts back to the client.
VS Server	<p>The VS Server service is responsible for various maintenance and monitoring related activities:</p> <ul style="list-style-type: none"> Watchdogs the RESPOND tasks in case of a task failure Creates new RESPOND tasks when a rule is created or modified by the user Provides a live image for the RESPOND task configuration wizard Provides the list of cameras for the Web Admin's camera activation dialog Creates the scheduled RESEARCH tasks Sends the outbound alerts to the outbound API and also sends alerts to the VMS clients that have real-time alerts integration (level 2a or above) Triggers the data maintenance activity Clears inactive sessions
Web Applications	BriefCam Web Services drive the main functionality for the web client and APIs. The service is stateless in nature, using the database and shared storage as its stateful backend. There are multiple web services provided by the component for both end users and administrators.

Server roles

There are seven core server roles that are required for a BriefCam multi-server solution:

Table 4. BriefCam server roles

Server type	Description	High availability
Web Services	Provide functionality required for the web client and developer APIs. These are stateless services that use the PostgreSQL database and shared storage to manage the application state.	Active/active stateless clustering that can have multiple active instances to scale out request demand and enable seamless failover using a NGINX load balancer.
Research Service Using Qlik	Optionally deployed in a multi-node architecture with shared access to a shared folder that contains application data.	Active/active. Failover should be configured based on Qlik's recommendations. A load balancer can be deployed to distributed load for seamless scale-out performance and high availability.
Processing Service	Also referred to as on-demand mode processing. This service is responsible for video decoding, rendering, object extraction, and classification that require one or more GPU cards.	Active/active. Multiple servers can be deployed at a single site to scale video processing requirements and provide high availability. All nodes in the solution are clustered. An application-level load balancing algorithm/policy determines the preferred node for a given session.
Alert Processing Service	Also referred to as real-time processing mode. This service supports the delivery of proactive responses to critical events for increased safety and security, with customizable alerts, alert reporting, and browser notifications.	Active/active. Multiple servers can be deployed at a single site to scale video processing requirements and provide high availability. All nodes in the solution are clustered. An application-level load balancing algorithm/policy determines the preferred node for a given session.
Fetching Service	The Fetching service is responsible for fetching the video footage from the VMS.	This service can be installed on multiple hosts to provide high availability.
Video Streaming Gateway Service (VS Server)	Used for camera management, VS service is responsible for starting, stopping, and managing additional BriefCam services	Active-Passive is the only deployment model for the VS Service. A second VS Service must be deployed on a different server and keep it on standby by making sure the service is stopped and not in use. The primary instance must be monitored to detect if it goes down. The standby instance must then be started either manually or automatically on the redundant server. Once the service starts, it connects to the database and continue where the failing service left off.
PostgreSQL Database Service	As BriefCam processes video, it detects and recognizes objects, along with information about their type and attributes. Objects are then classified according to different classes, attributes, color, dwell time, size, and speed. BriefCam processes the video stream once and then stores all metadata in the database.	PostgreSQL offers various ways to archive and replicate the primary database for backup, high availability, and load balancing scenarios. BriefCam recommends implementing a trigger-based master-standby replication.

High availability designs

An HA deployment requires additional virtual machines for redundancy, marked in the following diagram with yellow borders. A completely isolated HA installation for all critical services requires at least 15 virtual machines. In practice, more than one service can run on a single VM.

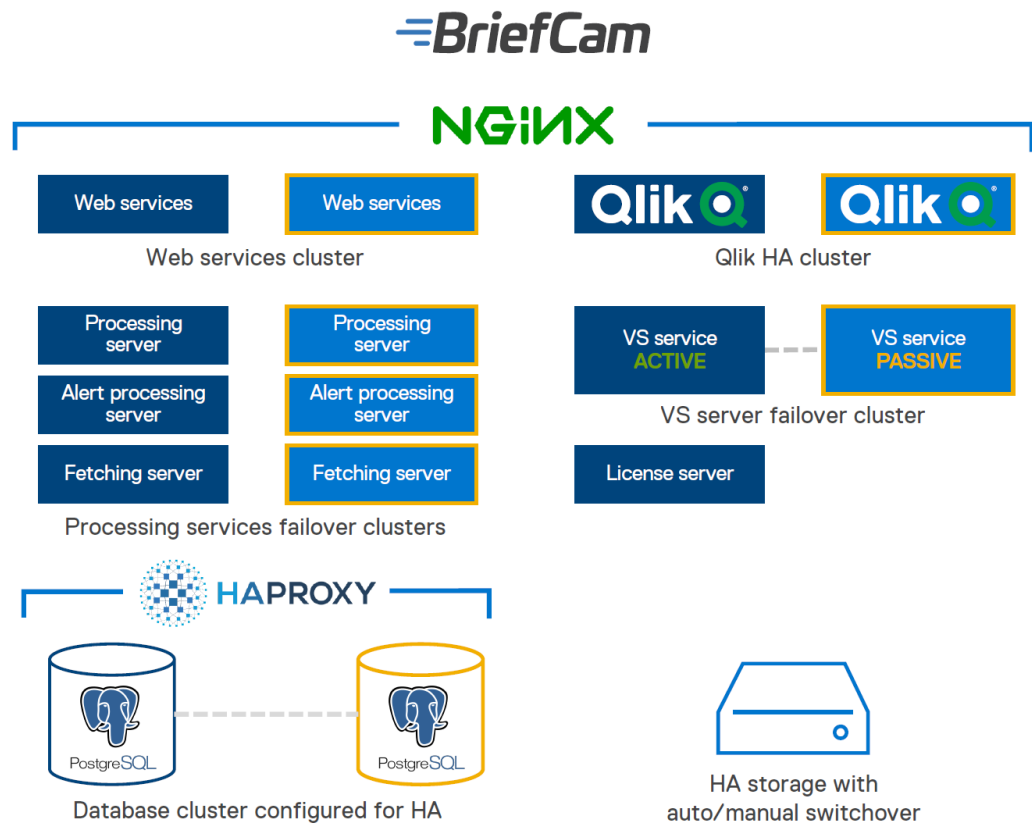


Figure 5. BriefCam HA architecture

High availability for microservices

The following BriefCam services were designed for HA, and therefore, do not require a third-party load balancer for clustering and redundancy. Their architecture natively supports one of the two redundancy models described in the following table:

Table 5. BriefCam services for HA

Service type	Services	HA redundancy model
Stateless	Rendering Face Recognition Fetching BI Rules Engine Alert Processing Maintenance Processing	These services support redundancy using a queue-based architecture. Service instances pull tasks directly from the database. There is no communication between services and each service that you add increases the system's capacity. When a single service fails, the system continues to work under reduced capacity.
Akka	LPR Matching BI Face Recognition Face Matching Filtering Lighthouse	These services use a service-mesh architecture provided by Akka.Net that automatically constructs failover clusters as more instances are added to the system. There is no need for load balancing or any additional configuration and each service that you add increases the system's capacity. When a single service fails, the system continues to work under reduced capacity. To construct a failover cluster, all you need to do is deploy multiple instances for each of these services.

VM designs

This section details the resources allocated to each VM in our validation lab. These configurations do not reflect supported configurations for an enterprise-class production environment. Sizing a solution depends on many environment-specific parameters and should always be reviewed with BriefCam Professional Services before implementation.

This following table details the VMs where the different roles will reside:

Table 6. VMs where roles reside

VM name	Cores	Memory	GPU	Storage	OS	Role
briefcam-db-1	16	64 GB	None	2 Disks c:\ 200GB d:\ 500GB	Win Server 2019	PostgreSQL Database
briefcam-ps-1 to breifcam-ps-3	16	140 GB	1 x A40 vGPU per VM	200GB	Win Server 2019	Alert Processing Services
breifcam-svr-4	16	128 GB	None	200GB	Win Server 2019	Web Services
briefcam-web-1	4	16GB	None	200GB	Win Server 2019	Management Consoles

High level architecture

The distribution of BriefCam VMs across the VxRail cluster was as follows:

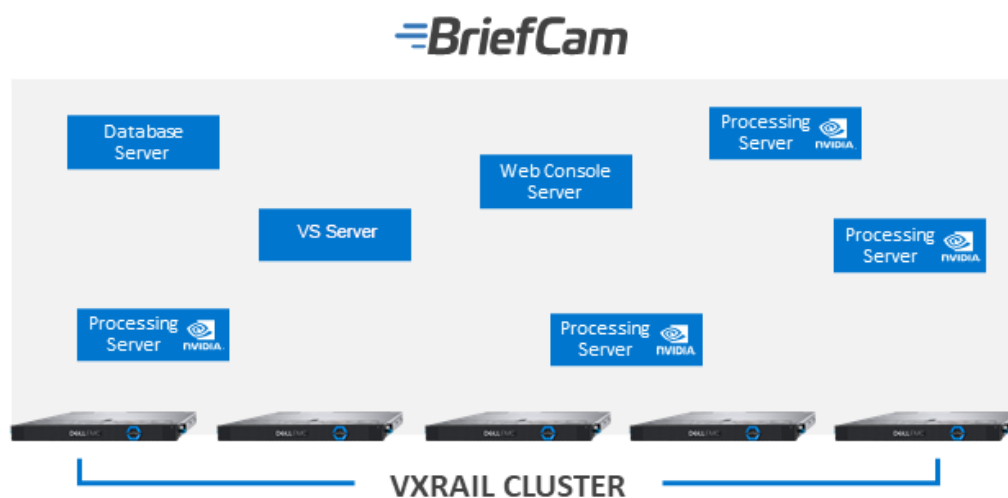


Figure 6. BriefCam VMs across the VxRail cluster

Ipsotek architecture

The Ipsotek VISuite Computer Vision platform extracts metadata from RTSP video streams allowing you to improve response times, drive operational efficiencies, and generate new revenue streams. This intelligent capability produces alarms for specific behaviors and events, allows operators to perform forensic searches, and generates dynamic dashboards.

Airports operate under strict government and aviation health safety guidelines which all staff must adhere to. VISuite supports multiple airport-specific workflows (use cases) and can automatically generate alerts for operators to manage and log. Examples include intrusion into aircraft-only zones, driving violations, and speeding. Alerts can be displayed on dynamic dashboards and updated according to the airport's operational requirements. VISuite can also be integrated with an email system to notify managers when certain events have occurred.

Scale-out architecture

The Ipsotek VISuite system architecture consists of Management Nodes, Processing Nodes, and Database Nodes that are interconnected as shown in the following drawing. There are also services that manage rules and user interfaces, perform Video Analytics, and provide event and metadata storage respectively. These roles could be deployed in a distributed manner and/or coexist on the same physical or virtual server.




Figure 7. Ipsotek scale-out architecture

Supported hardware redundancy

- N:N redundancy for active/active operation with zero downtime.
 - Requires hardware duplication.
- N+1 redundancy for Processing Nodes.
 - Requires load balancing.

NOTE: Redundancy does not require a software license.

Configuration best practices

The following guidelines are based on previous internal testing and reports from customer field experience. None of these guidelines were exceeded during testing.

- | | |
|-------------------------|--|
| Virtual Machines | <ul style="list-style-type: none">• Each VM is configured to the specifications provided by Ipsotek.• Each processing node is presented a full A40 card as a vGPU.• Uses a Maximum of two processing nodes per VxRail node. |
| CPU | <ul style="list-style-type: none">• Do not over allocate vCPUs.• The underlying CPUs are the faster 3 GHz model. |
| Storage | <ul style="list-style-type: none">• All Ipsotek nodes leverage the VxRail vSAN storage for operating system and auxiliary drives.• Data is protected with a RAID 1 storage policy to provide maximum performance and protection with the trade-off of reduced available capacity. |
| Networking | <ul style="list-style-type: none">• No network limits for Ipsotek. Sufficient bandwidth is needed at the processing nodes to consume the camera streams. |

High availability

It is critical to design a high availability deployment of Ipsotek to avoid loss of alerts during an unplanned outage.

HA options

It is possible to configure Ipsotek in two modes:

- N:N architecture
- N+1 architecture

Ipsotek N:N architecture

The N:N Architecture involves provisioning a full clone of all VMs in the Ipsotek system. This involves a full duplication of all servers in the cluster running on separate hardware. One cluster acts as the primary and both clusters have connectivity to the camera streams.

When running two identical clusters, Ipsotek has logic to ensure that configuration and alarms are synchronized to avoid duplication. This architecture should be used when full redundancy is required.

Ipsotek N+1 architecture

The Ipsotek N+1 Architecture provides high availability within a single cluster. In this configuration multiple management VMs are created and multiple Processing nodes are configured as a cluster.


With this approach, the processing nodes automatically join the cluster when they are booted. Once they join the cluster, the Processing nodes start to process camera streams. The camera streams are load balanced by the Management cluster to allow an even distribution of workload across the cluster.

This architecture was used for our testing.

Ipsotek roles

The Ipsotek system is made up of multiple roles and some of the critical roles for CV are:

Table 7. Ipsotek system roles

Role	Supports HA	Comments
Management Node	Yes	Central administration of the Ipsotek Cluster. Multiple nodes run in an Active/Passive manner using a Windows failover cluster.
Processing Node	Yes	Node where CV analytics is performed. Multiple processing nodes exist across the cluster. Workload is balanced across the nodes.
Database Node	Yes	Storage location for all Ipsotek config and metadata. Multiple nodes can be added together to form a cluster.  NOTE: A minimum of three data nodes is required to form a cluster. Only one was used for this testing.

The recommended approach for Ipsotek in a virtualized environment is to enable a GPU as a passthrough device. This means that it is not possible to easily migrate an Ipsotek processing node across the VxRail cluster. To combat this requirement, all four active nodes in the VxRail cluster must have a processing node configured. It is critical not to overload the four nodes and to allow enough capacity so that one of the nodes can become unavailable and the workload is still being processed. This is discussed further in the [High availability validation](#) section.

VM designs

This section details the VMs that the different roles reside on:

Table 8. VMs where roles reside

VM name	Cores	Memory	GPU	Storage	OS	Role
ipso-mgmt-1	8	34GB	None	200GB	Win 2019 Server	Ipsotek Management Node
ipso-mgmt-2	8	34GB	None	200GB	Win 2019 Server	Ipsotek Management Node

Table 8. VMs where roles reside (continued)

VM name	Cores	Memory	GPU	Storage	OS	Role
ipso-db-1	16	64GB	None	200GB	Win 2019 Server	Ipsotek Database Node
ipso-ps-1 to ipso-ps-4	8	64GB	1 x A40 per VM	2 Disks c:\ 200GB d:\ 400GB	Win 2019 Server	Ipsotek Processing Nodes

High level architecture

The distribution of Ipsotek VMs is managed by DRS, with the exception of the Processing nodes. The Processing Nodes cannot be migrated easily across the cluster due to the fact that they require passthrough GPUs.




Figure 8. Ipsotek VMs across the VxRail cluster

Validation

Topics:

- [Validation overview](#)
- [Performance validation](#)
- [Performance results and findings](#)
- [Performance summary](#)
- [High availability validation](#)
- [High availability results and findings](#)
- [High availability summary](#)

Validation overview

Our design validation used a 5-node VxRail cluster for application hosting, and a Dell PowerScale NAS storage device for accumulation and management of video data streamed to a Milestone Systems VMS. Based on recent sizing requests from customers, a common planning ratio for VMS archived streams recorded for historical analysis to "analytics" cameras requiring real-time alerting is approximately 5:1 or 20% real-time. The reason for this ratio is due to the fact that many cameras are static in nature and may not warrant the extra investment required to enable real-time analytics. We simulated a workload with approximately 30% analytics streams to test the upper bounds of typical processing requirements.

- We met our design goals for VMS performance by processing a total of 840 camera streams using 12 recorder VMs.
- We also met our design goal for BriefCam by performing real-time alert processing for 200 camera streams using 4 Alert Processing servers in an application mesh cluster.
- Ipsotek was also able to process 120 camera streams for real-time event alerting using 4 processing VMs.
- All high availability tests produced results consistent with the application and platform expected behavior.

Camera simulators

Due to the nature of testing in a lab environment, it is not always practical to add hundreds or thousands of real cameras to a system during testing. Video camera simulation software is often used for testing safety and security applications, including VMS and computer vision applications. Camera simulation allows system designers to stress test an entire end-to-end CCTV system while controlling complexity and cost. To enable testing we leveraged a camera simulator provided by Milestone Systems. The videos used were encoded as H.264 and wrapped with a Raw (not MP4) container. Enterprise-class video camera simulators allow testers to specify bandwidth requirements, quality, quantity, encoding, and source video files.

The camera simulation system used for testing our design needed to meet the following requirements:

- Camera simulators run on dedicated virtualized computing nodes independent of the VxRail platform.
- The proposed bit rate for all cameras is ~3.2 Mb/s.
- Specific source video files have been used to test the capability of the CV tool features we are targeting.
- Simulators connect and stream directly to the Milestone XProtect VMS using a dedicated network segment (VLAN).

Simulator hardware

All camera simulation was performed on a 2-node ESXi cluster using 2 x 750xa PowerEdge servers. We planned to simulate 840 cameras running on this simulation cluster. Each Simulator VM simulates 70 cameras with specific video files for Face Recognition.

Physical cameras

Our lab hosts a small number of physical cameras for performing selected specialized testing. For this Design validation, we need physical cameras connected to the Milestone Systems recorders that we used for high availability failover testing. Our physical cameras are capable of being redirected to a secondary IP address and port that is needed when a recorder fails whereas most camera simulators do have this feature.

Performance validation

The goal of our performance validation was to assess whether our proposed design could produce reliable planning input metrics for sizing a range of small, medium, and large deployment options. Most VMS and CV applications are sized by scaling out infrastructure based on per application service building blocks, e.g., video recorder service, database service, alert processing service, etc. We configured application service VMs to offer a good balance of performance and high availability. We monitored all key OS and application-specific performance metrics to see if we could detect any signs of resource exhaustion or process blocking. Our mixed workloads included processing a large simulated camera network with a Milestone XProtect VMS while simultaneously running BriefCam and Ipsotek real-time alert processing for a subset of the total camera streaming workload.

VM placement with DRS

A typical multi-vendor solution for a VMS/CV consists of a total VM count that is too difficult to manage with static rules and guidelines. VMware provides the DRS tool that will intelligently assess the layout of the application VMs on the VxRail solution to ensure even and reliable performance for the system. In addition to performance, it is also important to consider the impacts of high availability protection when defining DRS rules.

Placement constraints

The rules on placement of VMs are critical for small or medium size platforms due to the limited count of nodes in the system. This is less of a concern in a large system, but the rules must still be followed. Failure to follow the rules can result in some nodes that are over-allocated and not in a HA configuration.

The following application constraints were considered when configuring DRS for VM placement prior to the testing performed to produce this guide:

- A Maximum of 3 Milestone XProtect Recorders per VxRail Host
- A Maximum of 2 CV processing nodes per VxRail Host (assuming 2 GPUs available)
- Primary and secondary VMs for a system must not be on the same VxRail Host:
 - Ipsotek Management VMs cannot be on the same Host.
 - Ipsotek Database nodes cannot be on the same Host.
 - BriefCam Management nodes cannot be on the same Host.
 - BriefCam Database nodes cannot be on the same Host.
 - Milestone XProtect Management nodes cannot be on the same Host.
 - Milestone XProtect Database nodes cannot be on the same Host.
 - Milestone XProtect Hot and Cold Recorders cannot be on the same Host as the Primary recorders.

DRS Rule Details

1. No host or VM groups were needed for this design. The use of dynamic hosts that can get placed into maintenance mode results in static rules not working.
2. Briefcam- PS : A rule that ensures that all BriefCam processing nodes are kept on separate hosts
3. Mil Management: A rule that ensures that the two XProtect Management VMs are placed on separate hosts.
4. Ipso-ps : A placeholder rule for completeness when vGPU support is added to Ipsotek.
5. Max 3 Recorders per Host: One recorder from each group of three was chosen and forced to be placed on separate hosts.
6. Group Rec 1-3 : A rule that says **Keep Virtual Machines Together**. Creates a group that can run on any host.
7. Group Rec 4-6: see 6 above
8. Group Rec 7-9: see 6 above

9. Group Rec 10-12: see 6 above

The following screenshot shows the rules when configured with the vSphere web UI.

VM/Host Rules

Name	Type	Enabled
Brief Cam-PS	Separate Virtual Machines	Yes
Mil Management	Separate Virtual Machines	Yes
XDR-VDR-AntiAffiRule	Separate Virtual Machines	Yes
Ipso-ps	Separate Virtual Machines	Yes
Max 3 Recorders per Host	Separate Virtual Machines	Yes
Group Rec 1-3	Keep Virtual Machines Together	Yes
Group Rec 4-6	Keep Virtual Machines Together	Yes
Group Rec 7-9	Keep Virtual Machines Together	Yes
Group Rec 10-12	Keep Virtual Machines Together	Yes

VM/Host Rule Details

The listed 3 Virtual Machines must run on the same host.

Rule Members	Conflicts
mil-rec-2	0
mil-rec-1	0
mil-rec-3	0


Figure 9. DRS rules

VM placement for validation

After enabling VMWare DRS and starting all the VMs for the solution we observed the following distribution of VMs across the VxRail cluster:

Partner	Node 1	Node 2	Node 3	Node 4	Node 5
VxRail/vCenter		<ul style="list-style-type: none"> VxRail Manager VMware Center Server Appliance 			
Milestone XProtect	<ul style="list-style-type: none"> Backup Mgmt Server Recorder 1 Recorder 2 Recorder 3 	<ul style="list-style-type: none"> Recorder 4 Recorder 5 Recorder 6 	<ul style="list-style-type: none"> Main Mgmt Server Recorder 10 Recorder 11 Recorder 12 	<ul style="list-style-type: none"> Mgmt SQL Server Database 2 Recorder 7 Recorder 8 Recorder 9 	<ul style="list-style-type: none"> Mgmt SQL Server Database 2 Cold Standby 1 Cold Standby 2 Hot Standby 1 Hot Standby 2
BriefCam	<ul style="list-style-type: none"> Processing Server (with GPU) 	<ul style="list-style-type: none"> Processing Server (with GPU) 	<ul style="list-style-type: none"> Processing Server (with GPU) Web Server 	<ul style="list-style-type: none"> Processing Server (with GPU) 	<ul style="list-style-type: none"> Database Server VS Server

Partner	Node 1	Node 2	Node 3	Node 4	Node 5
Ipsotek	<ul style="list-style-type: none"> Processing Server (with GPU) 	<ul style="list-style-type: none"> Processing Server (with GPU) 	<ul style="list-style-type: none"> Processing Server (with GPU) 	<ul style="list-style-type: none"> Processing Server (with GPU) Ipsotek Management Server Datastore Server 	<ul style="list-style-type: none"> Processing Server (with GPU)

 **NOTE:** This is only one possible solution that could result from the application of DRS rules for placement.

Video workload

The testing workload was generated by 840 camera simulators using two different source videos. The details of the videos used are as follows:

Type	CV workload	Resolution	FPS	Bit rate
Busy Hallway	Face Recognition	1080	30	3.2 Mb/s
Isolated Stairwell	Restricted Zone	1080	20	2.8 Mb/s

Camera mapping to Milestone

The camera simulators ran on hardware outside of the VxRail cluster that hosted the VMS and CV processing services. The following table shows how the mapping of camera simulators to Milestone XProtect recorder VMs:

Camera Simulator	Milestone Recorder	Camera Count
Simulator-01	Recorder-01	70
Simulator-01	Recorder-02	70
Simulator-01	Recorder-03	70
Simulator-02	Recorder-04	70
Simulator-02	Recorder-05	70
Simulator-02	Recorder-06	70
Simulator-03	Recorder-07	70
Simulator-03	Recorder-08	70
Simulator-03	Recorder-09	70
Simulator-04	Recorder-10	70
Simulator-04	Recorder-11	70
Simulator-04	Recorder-12	70
		Total 840

Performance results and findings

Milestone XProtect performance tests

Single recorder performance test

VM specification

The VM specification for the single recorder test included:

- 6 vCPU
- 12 GB Memory
- 1 - 200 GB local vSAN disk (OS)
- 1 - 2 TB local vSAN disk (Tier 1 Storage)
- 1 - 10 Gb Network card

Test results

The performance data collected for a single Recorder processing a load of 70 camera streams is below.

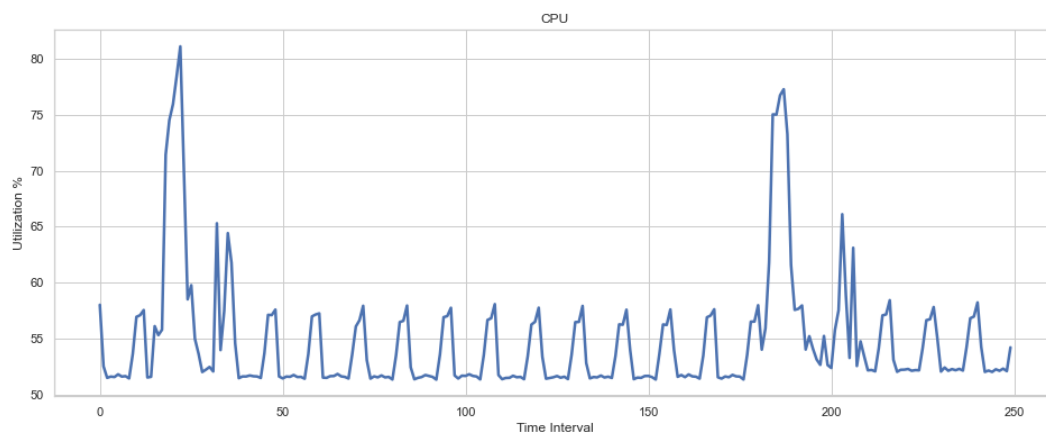


Figure 10. CPU results

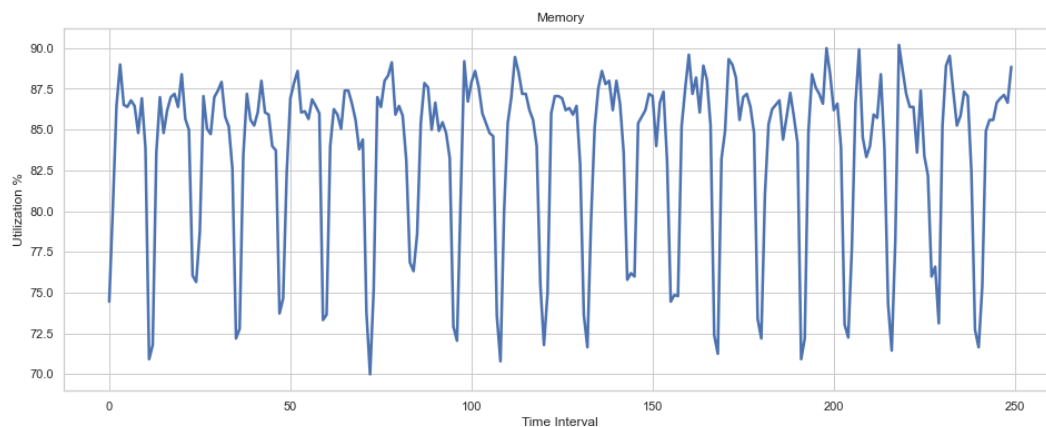


Figure 11. Memory results

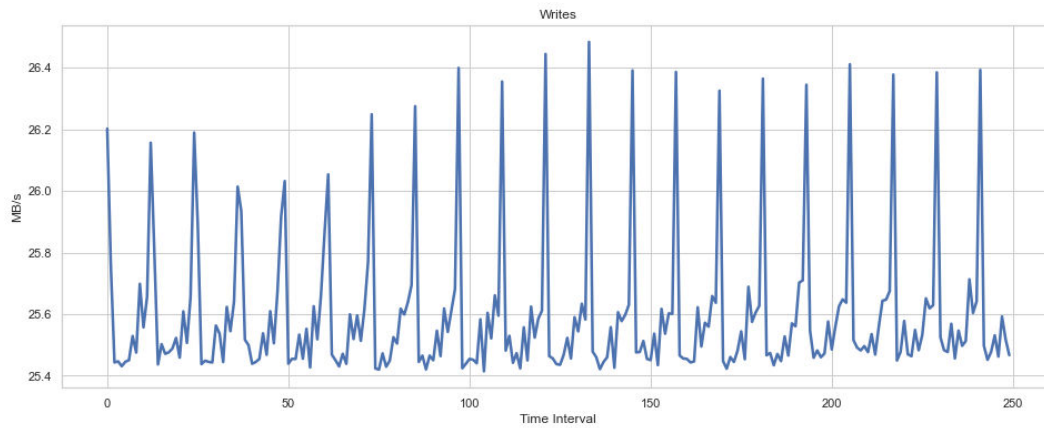


Figure 12. Writes results

Findings

- The percentage memory utilization for a recorder allocated 12GB and processing 70 camera streams was higher than expected with peaks around 90%.
- The maximum Tier 1 vSAN writes were ~26.5 MB/s. This is consistent with the best practices for Milestone sizing for a single VM where disk writes should be between 24 to 30 MB/s.

Overall

- The XProtect recorders performed well but we recommend that memory should be increased to 16GB when sizing a recorder.

Aggregate Recorder test for a VxRail node

Test specification

- Validate that 3 Milestone Recorders can run on a single VxRail host without exceeding the 96 MB/s thresholds set in previous Milestone testing by Dell.

We ran the following recorders on a single node of the VxRail cluster:

- Milestone Recorder 1
- Milestone Recorder 2
- Milestone Recorder 3

Test results

The disk write bandwidth data for a single VxRail host with 3 recorders streaming a total of 210 cameras is:

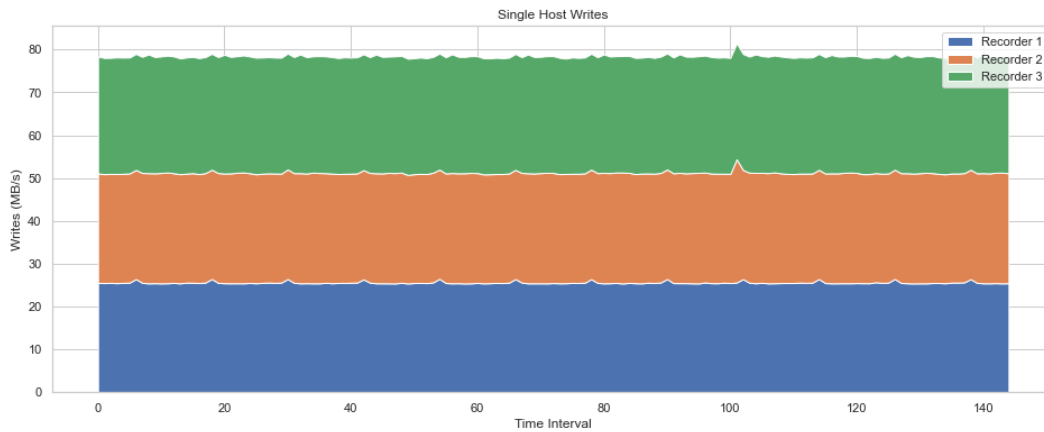


Figure 13. Single host writes

Findings

- The amount of vSAN writes was < 80 MB/s. This is below the 96 MB/s threshold established by Milestone Systems and Dell, and therefore, is not an issue.

Overall

- Hosting a maximum of 3 recorders per VxRail host is our recommendation for customer environments.

BriefCam performance tests

Our performance testing in support of the recommendation in this Design Guide was focused on the BriefCam RESPOND real-time analytics feature (Alert Processing Service). The integration between the RESPOND functionality and the Milestone Systems VMS uses a plug-in component on the BriefCam Alert Process server and the Real-Time Streaming Protocol (RTSP) interface supplied by Milestone Systems to access video streams with the lowest latency possible.

A real-time task processing request goes through several preprocessing steps before being able to generate alerts. The BriefCam RESPOND user interface allows operators to monitor the status of real-time tasks to track how many are queued, recovering (between the queued and processing state), and actively processing alerts. We tested performance by adding requests for new real-time alert processing tasks until we began to detect that the number of queued requests was increasing but were not able to change state and enter the processing stage.

Test specification

- Determine the maximum number of real-time alert processing tasks supported by a single VM with a full Nvidia A40 vGPU profile.
- Validate that the load balancing service in BriefCam will evenly distribute real-time alert processing tasks across a 4-node active/active cluster of BriefCam Alert Processing servers that are hosted on virtual machines each with a full Nvidia A40 vGPU.

Single VM test results

We tested both the maximum numbers of a single-use case workload stream (face recognition or person detection in a restricted zone) plus various mixtures of the workload specifications. In all tests, the results were consistent. In this test, cameras were added until we began to see that new streams were remaining in a queued state.

The total CPU and GPU Utilization in the build-up to processing for 50 Face Recognition Cameras are as follows:




Figure 14.

Detailed GPU utilization metrics are shown below:




Figure 15.

The Green line in the above graph shows that the BriefCam Alert Processing Service does not perform any re-encoding of the video format that is streamed from the Milestone Systems VMS. The BriefCam Milestone Systems plug-in installed on the Processing Server knows the format of video produced by the RTSP service (H264) which is also known to the processing engine. This VMS-specific product integration reduces GPU workload and allows more resources to be used for alert processing.

Multi-server load balanced Real-time Alerting

The BriefCam Alert Processing Service can function as a scale-out active/active cluster by installing and configuring multiple instances of the service that run on different servers. When multiple service instances are available, a cluster mesh service detects how many service instances are running and on what machines. When new real-time alert processing tasks are requested, the request is queued using the PostgreSQL database. The service mesh then assigns tasks from the queue to available servers that are running the alert processing service.

We configured 2 VMs running the BriefCam server components to process real-time alerts. Our goal was to validate that the total number of streams processed would be allocated evenly among the cluster nodes. The maximum number of streams processed during the cluster test was 80, There was no request queuing at that level of workload consistent with the finding that the single server maximum is 50 streams shown above.

We added new camera streams to the workload in groups of 10 with a several-minute pause in between to verify that all streams were in processing mode before proceeding. The table below shows the comparison of resource utilization for the two load-balanced Alert Processing servers.

Table 9. BC63-SVR-3

	CPU % utilization			GPU % utilization			GPU video decoding % utilization		
	Min	Avg	Max	Min	Avg	Max	Min	Avg	Max
BC63-SVR-3	66	75	88	75	81	84	13	37	71
BC63-SVR-4	62	75	88	60	71	78	18	45	80

Findings

- The GPU is the most heavily used resource for Alert Processing servers.
- The BriefCam active/active load-balancing is deficient in spreading the workload across the servers in a cluster.

Ipsotek performance tests

The Ipsotek alert processing testing contained a mixture of Face Recognition and Restricted Zone camera simulation scenarios. The Ipsotek platform supports multiple tracking modes that can be configured for a camera:

- Video Analytics** Traditional pixel-based analytics.
- AI - Normal** AI based tracking of full people. This was used for Restricted Zone tests.
- AI - Crowded** AI based tracking for crowded scenes where head and shoulders are tracked.
- Face Detection** Focus on detecting faces. This was used for Face Recognition tests.

A "Face in Watchlist" rule was created to generate a report for people identified from a camera stream and on a watchlist. We used this watch list scenario for all Face Detection camera processing. In addition, an "In Zone" rule was created to detect people entering a restricted zone.

Ipsotek uses the Nvidia GPU encoder to reencode the video received from different cameras and generate synchronized video with I-Frames produced every second at a fixed bit rate. This allows Ipsotek to:

- Estimate storage requirements based on the retention policy.
- Access video and images on disk to an exact second without needing to seek in mp4 segments.
- Stream video over HLS protocol with a minimum buffering time of 2 s to be viewed on a web interface.
- Ensure that we get a low latency user experience when an operator is requesting to random access images and video on disk.

Test specifications

- Determine the maximum number of real-time alert processing tasks that can be supported by a single VM with a full Nvidia A40 vGPU.
- Determine the maximum number of real-time alert processing tasks that can be supported by a 3-node active/active cluster of Ipsotek Processing servers hosted on virtual machines each with a full Nvidia A40 vGPU.
- Use a combination of 45 face recognition and 45 restricted area/person detection workloads and collect performance metrics. In particular, focus on GPU Encoder and Decoder usage.

Single VM test results

We tested both the maximum numbers of a single-use case workload stream (face recognition vs person detection in a restricted zone) plus various mixtures of the workload specifications. In all tests, the results were consistent. Cameras were added until the max of 30 cameras were active and no more could be processed.

The total CPU and GPU Utilization in the build-up to processing for 30 Face Recognition Cameras are as follows:




Figure 16. CPU and GPU utilization for face recognition cameras

The breakdown of GPU utilization is as follows:

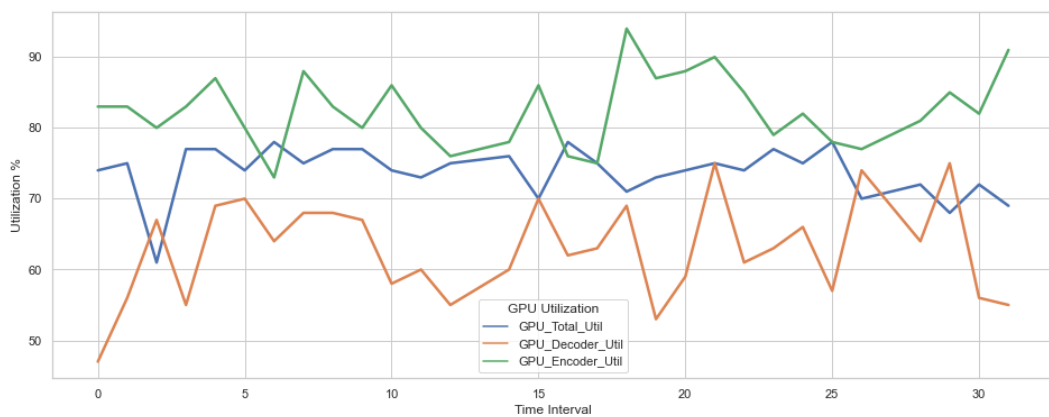


Figure 17. GPU utilization

Findings

- The max number of Face Recognition or Restricted Zone cameras is consistent at 30 cameras per Ipsotek processing server.
- The encoder component of the A40 GPU was exhausted during this test.

VM cluster test results

The Ipsotek platform does not currently support vGPUs, so all GPU resources must be assigned to the VM as a passthrough device. This means that it is not possible to migrate VMs as was shown in High availability validation. To build a viable Ipsotek cluster in a virtualized environment, it is necessary to provision all Ipsotek Processing nodes at system setup time. The goal is then to load the system so that sufficient capacity exists across the cluster to ensure that during a failover the cameras can migrate to an available node and continue processing.

In this testing, the approach was to load four of the five nodes in the cluster to their maximum. This ensures that there is sufficient capacity to disperse the workload across the cluster in a failure scenario. A total of 60 Face Recognition Cameras and 60 Restricted zone cameras were enabled.

The maximum count of active cameras across three processing nodes was 120.

Findings

- The maximum number of real-time alert processing tasks was 30 for a single node. We also validated that this result was consistent across the three VxRail nodes that were allocated for CV application hosting.
- The CPU was ~50% when all cameras were processing so 8 CPUs per Ipsotek processing node is sufficient.

- When 30 cameras were enabled on a single node, the GPU utilization was at ~50% but the GPU Hardware Encoder approached 100% and no further-incoming video streams were accepted showing that video encoding is the most resource-intensive component of the alter processing workload.

Full system performance tests

The purpose of this test was to validate if a combination of three applications from three different vendors can share a common platform using VxRail and VMware without introducing processing delays. We have the individual application results from above for our baseline. Our performance data for this test was collected while the following workloads were processed in parallel.

- A total of 840 cameras are being recorded in Milestone XProtect.
- 200 BriefCam real-time alert processing tasks analyzing camera streams from the Milestone Real-Time Streaming Protocol (RTSP).
- 120 Ipsotek Face Recognition and Restricted Zone tasks analyzing camera streams from Milestone XProtect using RTSP.

Overall cluster performance results

The following chart shows a snapshot of the CPU utilization of the DRS cluster at a point in time: during the full system testing.




Figure 18. CPU utilization of the DRS cluster

Findings

- The workload is distributed across all 5 nodes in the cluster using DRS.
- No issues when streaming 840 cameras.
- Sufficient capacity exists that allows one node be taken offline for maintenance.
- Sufficient capacity exists to handle any spike in the load.

Overall

- This cluster is able to handle the workload without issue.

VM level view

When the full system load test with 840 cameras was running, the top 10 most active VMs were identified by CPU MHz consumption.

The following charts shows those top 10 busiest VMs in the cluster:




Figure 19. The cluster's top 10 busiest VMs

Findings

- The Milestone XProtect recorders are consuming the most CPU across the system.

Overall

- It is expected that the Milestone Recorders are utilizing about 10 MHz per VM. The VMs are allocated 6 vCPUs and are running at less than 75%.

Performance summary

The full system performance test with all 840 cameras storing video and running analytics on 320 cameras performed in line with expectations. No part of the system was over loaded.

During testing it was observed that there is an improvement between BriefCam 6.1 and 6.3. Previous test results showed 30 analytic camera streams per A40 whereas this testing shows 50 cameras per A40 is supported with BriefCam.

High availability validation

As described in the [Solution Architecture](#), a key element of this design is to provide high availability for every application. Since VxRail, VMware, and our partner solutions all offer features for implementing high availability, there are many possible implementation design patterns. We present only one of the many possible approaches to HA in this section but there are many more options that can be explored based on the specific needs of your environment.

The following high availability features were tested against our design:

- VxRail high availability
- VMware vSphere vMotion (for the BriefCam VS Server service)
- VMS and CV applications high availability

VxRail high availability

VxRail provides a seamless, curated, and optimized hyper-converged experience based on the joint engineering efforts of Dell and VMware. This deep integration combined with the simplicity of VxRail provides an ideal platform for implementation across most core, edge, and cloud environments. The VxRail HCI System Software coupled with the performance of next-generation PowerEdge servers offers an industry-proven scale-out high-performance platform. Our CV for Transportation platform design uses a 5-node VxRail cluster that was validated through the execution of simulated workloads. For detail on how Dell VxRail can deliver 6x9s High Availability see our [Product Brief](#).

Node planned downtime validation

Description	This test simulates the scenario where a hardware node of the VxRail cluster must be taken offline for maintenance.
Steps	<ul style="list-style-type: none">• The Ipsotek application VM running on the affected node uses a GPU in passthrough mode and therefore must be manually shut down since it cannot be migrated to another node in the cluster.• An administrator changes the node state to "Maintenance Mode" using the vCenter console.• Migrate any eligible workloads off the affected host.• Perform maintenance on the affected host, such as patch and reboot.• Exit maintenance mode and return the VxRail cluster to normal operation.
Expected results	Any workload that can migrate is redistributed across the cluster. The Milestone Systems Management Server can go offline without impacting any configured and running services. No new cameras or services can be configured while the Milestone Management server is offline. All services that cannot migrate but are protected by application-level high availability continue to provide services with zero downtime.

Results

- After the host was successfully put into Maintenance Mode using the vCenter console, all active VMs were migrated to other hosts in the cluster with capacity.
- The host that was shut down contained:
 - 3 x Milestone Systems Recorders
 - 1 x Milestone Management SQL Server Database
 - 1 x BriefCam Processing Node
 - 1 x BriefCam Web Service
 - 1 x VxRail Manager
 - 1 x VMware Center Server Appliance
 - 1 x Ipsotek Management Server
 - 1 x Ipsotek Datastore Server
 - 1 x Ipsotek Processing node (Not migrated)
- All cameras connected to Milestone recorders continued processing as normal.
- Zero Events on the BriefCam Admin Console. All cameras continued processing as normal.
- Zero alerts on the Ipsotek Management console. All cameras continued processing as normal on other servers.

Unplanned node downtime validation

In the unlikely event that there is an unplanned outage of a VxRail node, it is important that the system continues to operate with minimal disruption to services.

Description	This test simulates a hardware failure of a node in the VxRail cluster.
Steps	<ul style="list-style-type: none">• Simulate a failure by forcing the shutdown of a selected VxRail host.• Use the Dell iDRAC interface to log in and forcibly shut down one host at the hardware level.

	<ul style="list-style-type: none"> This test procedure ensures that the system has no opportunity to offload any workload to another host before the shutdown occurs.
Expected results	<p>All VMs that are residing on the failed node should be migrated automatically to an available location on the cluster.</p> <p>This migration process will take a length of time determined by the number and type of VMs running on the host that fails. System operators must configure the application HA when available to minimize any loss of service.</p>

Results

- One VxRail node was forcibly shut down using iDRAC.
- vCenter displayed an error showing the host offline .
 - All VMs showed their status as "Disconnected".
 - VMs were not operational at this time.
- The VxRail HCI system auto migrated the VMs assigned to the failed host to another host with capacity in the cluster except for the Ipsotek VM
- After migration, all VMs except the Ipsotek server were powered on automatically.
- The Ipsotek VM was not able to migrate since it does not support vGPU. This VM stayed in a down state with the error "vSphere HA virtual machine failover failed".
 - When the original host (Host-02) became available the Ipsotek VM was automatically rebooted and continued processing as expected.
- All migrated VMs were rebooted after 2 minutes and were processing workload a minimum of 5 minutes after the failure.
- The specific VMs that were part of this test were:
 - 1 x Milestone Main Management Server
 - 1 X BriefCam Processing Node
 - No issues as part of a Cluster. See BriefCam validation below.
 - 3 x Milestone Recorders
 - These recorders stopped processing for a minimum of 5 minutes.
 - We also test hot and cold backup recorders that can be configured for key servers in the next section.
 - 1 X Ipsotek Processing node
 - This cannot migrate. It remained down for the duration of the outage.
 - This is expected behavior due to no vGPU support and has been built into the [Ipsotek Architecture](#).

High availability results and findings

All three applications that were part of the validation were tested for HA functionality:

- Milestone 2022
- BriefCam 6.3
- Ipsotek 11.7.1

Each application was set up following the vendor recommendations. Additional HA validations were performed in addition to the VxRail HA tests.

BriefCam HA results

Protecting mission-critical CV applications is becoming an increasingly important aspect of facilities management. BriefCam supports a rich set of active/active as well as active/passive HA features. In this validation, we tested the native active/active HA features of the BriefCam alert processing server and active/passive protection of the VS Server. One of our design goals is to provide the best combination of native and VMware HA to protect everything from a specific VM to the full cluster to avoid operational disruption.

The BriefCam application has full support for the following features when building an HA solution.

- Redundant Microservice architecture
- Backup Management console

- Backup Database servers

Description	This test simulates a failure of a machine that runs the real-time Alert Processing Service.
Steps	<ul style="list-style-type: none"> • Select any BriefCam Alert processing VM and shut it down. It must have active camera traffic. • Monitor camera behavior in the Web console. • Measure the time taken for the camera to be moved to an active node in the cluster.
Expected results	<p>It is expected that when an Alert Processing Service goes offline other nodes running the Alert Processing Service will take over processing and alerts generation.</p> <p>The time taken for the processing to switch from a failed node to an active one will be measured.</p>

Results

- When the Alert Processing VM was shut down, the cameras that were running on that server went into the status "Active (Queued)"
- No alerts were being triggered while in this state.
- After a minimum of 5 minutes, the camera streams were picked off the database queue and launched on another processing node in the cluster.

Ipsotek HA results

Ipsotek supports additional application-level security on top of the VxRail provided HA. This security protects against an issue with a specific VM when the full cluster is still operational.

The Ipsotek application has full support for building an HA solution. It is possible to build two fully redundant systems processing systems and also to build redundancy in to a single cluster.

The single cluster option was used for this testing and has the following features:


- Multiple active Camera Alert processing nodes
- Backup Management console (active/passive)
- Multiple Database servers (min 3-node cluster)

The following scenario was validated during this phase of testing.

Description	This test simulates a failure of a machine that runs the core Camera Alert processing.
Steps	<ul style="list-style-type: none"> • Select any Ipsotek processing VM and shut it down to simulate failure. • Monitor camera behavior in the Ipsotek configuration tool. • Measure the time taken for the camera to be moved to an active node in the cluster.
Expected results	<p>When a Camera Processing VM is shut down other nodes take over and process the camera alerts instead.</p> <p>The time taken for the processing to switch from a failed node to an active one is measured.</p>

Results

- Once the VM was shut down, 15 cameras went in to "Active - No Video" mode in the Management console.
- The cameras were automatically distributed to among the remaining three active Ipsotek Processing nodes.
- After a maximum of 30s, all cameras were back processing alerts. A camera is reassigned to a new cluster member after 20s downtime. The Management console refreshes the camera status every 30s.
- The "Processing Node Information" view on the Management console shows that one node is not part of the cluster and that the workload has been distributed among the remaining active nodes.
- Once the Processing node was restarted, it automatically rejoined the cluster.

 **NOTE:** Once an Ipsotek Processing node joins a cluster, the cluster will not automatically rebalance to adjust the workload. If this is required cameras can be disabled and enabled to distribute to the new nodes.

Milestone HA results

High Availability is an important consideration for security systems deployment to avoid loss of camera stream data during an unplanned outage. For larger environments, Milestone Systems recommends hosting SQL Server on a dedicated server to provide adequate latency for many devices and/or many event transactions. For small to medium environments all four of the following VMS components can run on a single computer:

- Management server
- Event server
- Log server
- SQL Server

When using dedicated SQL Server installations, the HA design for the three remaining services (management, events, and logging) must be planned for separately. During our validation we installed two SQL Server database servers in an Always-On Availability Group for high availability. In environments where the uptime needs of the other three services is critical, consult with Milestone Systems and your system integrator to understand the options available.

SQL database failover

Our test began with two Microsoft SQL Servers configured with AAG and all Milestone databases restored to the second server and log replication active. We also had two Milestone Management Servers installed and connected to the Virtual Server name of the AAG.

- The test began by failing the databases to become active on the second database server in the AAG (failed from mil-db-1 to mil-db-2).
- All camera streams contained to be recorded without interruption.
- The management server UI on mil-dir-1 was used to disable 2 camera streams.
- A login was made to mil-dir-2 and the two cameras impacted in step 3 were in the disabled state.
- The two cameras were then successfully enabled on mil-dir-2
- This test confirms that SQL Server can successfully be protected against a single VM failure using SQL Server AAG.

Recorder hot standby failover

Hot standby failover recording server setup requires a dedicated failover recording server for each protected recording server. This one-to-one mapping allows the system to quickly transition a failover recording server from "standby" mode by synchronizing the correct/current configuration of the recording server it is dedicated to.

We began the test with two VMs, the primary recorder (mil-rec-1) and one hot backup recorder (mil-hot-1).

- The data folder configured for storing video recording locally is empty on mil-hot-1, the hot standby server.
- vSphere was used to shut down the primary recorder (mil-rec-1) that is then shown as offline in the Management Server.
- We observed data being written to local files on the hot standby server.
- The data from the local file storage is then merged back to the primary server when fallback is initiated
- Checking the video archive for a camera attached to the primary server shows approximately 15 seconds of missing recording when the initial failover occurred and 7 seconds of missing video when the fallback occurred.

We also checked the potential impact of a CV application processing camera streams mapped to the primary server before performing a hot standby failover test. The charts below show the processing impact on a BriefCam alert processing server analyzing several camera streams as the processing moves from the primary recorder to the hot standby and back again. We first see the primary failover around time interval 20 and fallback around time interval 50. The BriefCam user portal showed that alerts were continuously processing through both events even though there were losses of 15 seconds and 7 seconds of video data during failover and fallback respectively.




Figure 20. Hot failover CPU versus GPU




Figure 21. Hot failover GPU details

Recorder cold standby failover

We tested several scenarios using the cold standby features for Milestone XProtect video recorders. Our setup started by creating 3 VMs and installing the Systems Failover Server service and the Failover Recording Server service described in the Milestone Architecture section of this Design Guide.

We then created two Failover Server groups called cold-failover and cold-failover-2. The mapping of server VMs to Failover Groups was:

Failover group	Servers
cold-failover	mil-cold-1 mil-cold-2
cold-failover-2	mil-hot-12

NOTE: We apologize for any confusion caused by the naming of a cold standby server "mil-hot-12". The VM was repurposed from its original intended use.

We completed our initial setup by configuring the cold-failover group as the Primary Failover Server Group and the cold-failover-2 group as the Secondary Failover Server Group for the mil-rec-3 video recorder.

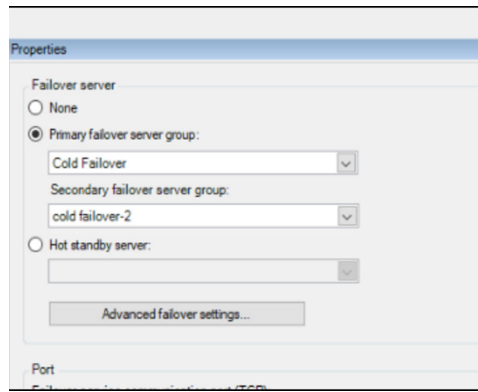


Figure 22. Cold-failover properties

The expected behavior with this configuration is that when mil-rec-3 fails or goes offline recording for any attached cameras will be transferred to either one of the two servers in the Primary failover server group. Failover server groups can protect more than one recording server so that in the event that both servers in the primary group are in service from prior failures when mil-rec-3 fails, it will then failover to a server in the Secondary failover server group.

Our first test was to create a failure on mil-rec-3 with all three cold standby servers in the two groups available. We had 2 physical cameras streaming to mil-rec-3 receiving approximately 15mbps of video data. We then powered down mil-rec-3 simulating a failover and the camera streams were transferred to mil-cold-1. There were approximately 37 seconds of missing video data that was not recorded during the failover. After the first simulated failure, we still had two additional cold standby servers available.

For our second test, we simulated a failure for mil-cold-1 and we saw nearly identical behavior. The two camera streams were transferred to mil-cold-2, the second failover server in the Primary failover server group, and observed about a 38 seconds gap in the video archive for those two cameras.

Our final test was to simulate a failover of mil-cold-2. We saw the same results as the two previous tests, but, this time the recording was transferred to the single VM in the Secondary failover group - mil-hot-12. The amount of lost video data this time was only about 26 seconds. This may be the result of this being the last available server and therefore there is no logic for choosing which server to failover to but we cannot verify that.

We also configured real-time alert processing for the two physical cameras through integration with our BriefCam CV application. The video from the physical cameras was of extremely low quality for use in a CV application. We were able to occasionally detect a person in the video using a people counting algorithm. The two charts below show resource consumption on the BriefCam Alert Processing server configured for this testing. There was minimal impact on resource utilization during the failovers.




Figure 23. Cold-failover CPU versus GPU




Figure 24. Cold-failover GPU details

We did see some interruptions in the alert rule processing in the BriefCam event log. One of the two streams would change from the Active Processing state to Recovering, then into the Processing Queue, and finally resume Active Processing. The duration of the disruption was roughly equivalent to the duration of the last video recording. There was no manual intervention required on the BriefCam application to resume stream processing during any of these standby failover tests.

High availability summary

The VxRail HCI system with VMware virtualization plus Nvidia A40 server-class GPUs. NVIDIA vGPU compatibility with the Dell Validated Design for AI platform provides important high availability options when implementing large systems with many virtual machines that augment the features of our application partners.

The following high availability features were tested against our design:

- VxRail high availability
- VMware vSphere vMotion (for the BriefCam VS Server service)
- VMS and CV applications high availability

The high availability features of VxRail systems and VMware virtualization are design components that are already familiar to many IT and OT (operational technology) professionals. These features provide additional capability for managing high availability that can complement the native features of the CV and VMS applications.

All VMS and Computer Vision products were installed on a single VxRail 5-node cluster and tested to demonstrate the value of high availability operations on a single common platform. The VMS high availability testing was performed with physical cameras since our simulation software did not support dynamic rerouting of output streams. We also tested with simulated camera streams being archived to the VMS simultaneously with two CV applications performing real-time analytics on a subset of the simulated camera streams for our CV applications HA validation.

This section confirms that the VxRail with VMware virtualization and NVIDIA GPUs provides a reliable Dell Validated Design for AI platform for hosting VMS and CV applications from different vendors on a common platform with a complete set of enterprise-class high availability features.

Sizing the solution

Topics:

- [Overview](#)
- [Sizing guidelines](#)
- [Scaling guidelines](#)

Overview

There are many environmental and use case requirements that will impact the sizing of a VMS plus CV application solution. Many of these factors cannot be simulated in a lab environment, however, with the right source videos and camera simulators a relatively large implementation can be modeled closely enough to gather useful sizing data.

Our design evaluation was targeted at collecting data for three sizes of transportation systems based on the total number of cameras and the number of cameras that require processing with a CV application, primarily for alert generation. The table below shows our three category definitions.

Table 10. Transportation system category definitions

Size	No. of cameras	No. of analytic cameras
Small	250–500	Up to 125
Medium	501–1000	Up to 250
Large	>1000	> 250

Our sizing categories above reflect the reality that most customer systems will need a mixture of cameras with some near real-time analytics as well as those that are intended for archive only in the event there is a need for historical video for investigations. In this testing, we are simulating a ratio of 1 analytic camera for every 5 total cameras installed. This will vary across customers and workload needs but 5:1 is the baseline used for all our sizing data collection.

The design can be applicable for environments with fewer than 250 cameras (small) but should also be useful for system designers working with airports with greater than 1,000 cameras (large).

Sizing guidelines

When sizing a VxRail cluster for implementation of our solution, it is important to first allocate the capacity that is needed to run any management VMs for the VMS and CV applications. The following management node types should be carefully designed for the proper scale and HA requirements of the deployment.

- | | |
|---------------------------|---|
| Milestone XProtect | <ul style="list-style-type: none"> • Redundant SQL Server Databases • Redundant Directory Servers (optional) • Backup Archivers (optional) |
| BriefCam | <ul style="list-style-type: none"> • Redundant Web Services Consoles • Redundant Databases |
| Ipsotek | <ul style="list-style-type: none"> • Redundant Web consoles • Redundant Databases |

This design does not include recommendations for every VM and service that could add benefit from an HA implementation in a mission-critical environment. Our design should give a reasonable starting configuration for each type of VM. We note the sizing specifics for each application in the following sections.

System sizing

Each VxRail node has the same configuration as described in the [Physical architecture](#) section.

Storage sizing

For mission-critical safety and security solutions, it is necessary to have sufficient storage to archive video content for analyzing historical events. The retention period varies per customer and can be multiple PB in scale for large environments.

The technology used for VMS recording during this testing was the PowerScale platform. We have previously validated with the XProtect VMS and found that it provides flexible, scale-out file storage with the performance needed by large-scale camera networks. PowerScale is easily configured with XProtect for Tier 2 storage.

For more information about Dell Storage with Milestone, see [Sizing Guide-Dell EMC storage with Milestone XProtect](#).

Cameras per node

Each VxRail node in the cluster is capable of hosting up to three Milestone XProtect primary recorders processing 70 camera streams each. We were also limited to a maximum of one Ipsotek processing server and one BriefCam Alert Processing VM per node, each with one A40 GPU allocated.

- We tested with 840 camera simulator streams at an average of 3.2 Mb/s with 12 primary recorders.
- We tested processing 30 analytic camera streams with Ipsotek and 50 streams for BriefCam per node used for Face Recognition and Restricted Zone generating real-time alerts.

As shown in the Performance validation section we tested running 210 cameras across three recorders, the total server bandwidth consumed is less than 640 Mb/s.

Similarly, if processing for fewer analytics cameras is needed then it is possible to remove one or more Nvidia A40 GPUs from the VxRail node.

Scaling guidelines

The same core platform design can be used for small or large systems. The following sections list some of the differences.

Small and medium system scaling

The platform design architecture we tested is ready for use with a medium airport system (500 to 1000 Cameras).

It is important to note that at a minimum a single instance of all the management VMs needs to be provisioned even with a small cluster.

Other options to consider are:

- Deploy a minimum of 4 active + 1 standby VxRail nodes.
- If fewer analytic cameras are needed, reduce the GPUs to one per node. This is preferred to going with less capable cards and gives the best option for future expansion.
- Memory cannot be reduced for the nodes since the VMS or CV management VMs can require large memory and need to fit in the same VxRail cluster.
- While it might be possible to reduce the CPU core counts for very small systems, it is recommended that validation be performed before finalizing the specification.

With this approach, it is possible to scale a Small or Medium platform by adding up to two A40 GPUs per host, then adding additional hosts to the cluster if the environment needs expand over time.

Large system scaling

A large system has some notable characteristic differences. For example, the management VMs carry less overall system overhead when distributed across a large cluster.

The following guidelines can be used when scaling up to a large system:

- Calculate the number of nodes required based on:
 - Each node can have a max of 1440 Mb/s or 420 Cameras @3.2 Mb/s.
 - Each node can host two CV applications processing a max of 80 analytic cameras total.
- Allocate two additional nodes for all management VMs (such as Primary or Secondary databases)
- Add one or more standby nodes for HA operations.

Example: Building a 5000 camera system

Our system requirements consist of:

- Enough capacity to process 5000 cameras @3.2 Mb/s
- Resources capable of running CV applications processing video from 1000 real-time analytic cameras

Number of processing nodes	Camera archive = $5000/420 = 12$ nodes Camera analytics = $1000/80 = 13$ nodes Number required = the maximum from above = 13 nodes
Number of management nodes	2 nodes
HA capacity	2 nodes
Total	17 nodes

Backup and restore operations

The Computer Vision platform resides on vSAN across a 5-node cluster with one dedicated fault-tolerant node. So data loss occurs only if two nodes are lost simultaneously or the second node fails before first node becomes available after fault. The backup and restore operations detailed here relate to application data backup and restore.

Topics:

- [BriefCam](#)
- [Ipsotek](#)
- [Milestone Systems](#)

BriefCam

There are two schedulable backup tasks that run daily for BriefCam environments:

1. PostgreSQL Database Backup
2. Research Dashboards Backup

Two separate Windows Task Scheduler Jobs are created during installation with default times of midnight for the database backup task and 2 AM for the Research Dashboards backup task. The scheduled backup times can be changed in the Windows Task Scheduler.

PostgreSQL database backups

Restoring the PostgreSQL database from a backup requires a service outage by stopping the VS Server service for the deployment. It is highly recommended that administrators contact the BriefCam Support team before performing a database restore.

Depending on the circumstances leading up to the need to restore a database, you may want to restore into an existing database or create a new database from a backup. The *BriefCam Administrators Guide* has script examples for both types of restoration.

Restoring database backups of research dashboards (Qlik)

Research data backups are written to a folder that is configurable after installation by editing the `BI_Backup` batch script and changing the value of the `BACKUP_PATH` parameter. It is recommended to set up a second backup task with a folder on another machine or copy the primary backups to a different machine for redundancy.

Restoring backups of the Research Dashboards uses the Research server installer to create a clean environment on a new server where the data from the previously failed server is transferred. The installer creates a fresh database when the install repository database service component is started. The administrator then uses a script to drop the fresh database, create a new database and restore the backed-up database from files. The full procedure is described in the *BriefCam Administrators Guide* documentation.

Ipsotek

The Ipsotek VConfigure tool has support for system Backup and restore. This option is visible under the **Server Settings** tab. It is highly recommended to backup the application configuration at regular intervals. The backup is an XML file and can be inspected as needed.

If there is a need to restore, the administrator can go to the previous backup and select restore.

During the restore process, there are options to restore some or all of the configuration. Some examples are to restore:

- Cameras
- Presets
- Actions
- Rule Assignments

Throughout the Ipsotek VConfigure admin tool, options exist to export individual items also. For example, individual rules can be exported and shared. This could be a useful feature if rules are developed in a nonproduction environment and these rules must be pushed to a production scenario.

Milestone Systems

Milestone recommends that you make regular backups of your system configuration as a disaster recovery measure. While it is rare to lose your configuration, it can happen under unfortunate circumstances. It is important that you protect your backups, either through technical or organizational measures.

Backing up your system configuration can take a significant amount of time, however, your system will stay online during backup. The backup duration may depend on:

- Your system configuration
- Your hardware
- Whether you have installed the SQL Server, Event Server component, and the Management Server component on a single server or several servers

All relevant system configuration files are combined into one single .cnf file that is saved at a specified location. During the backup, all backup files are first exported to a temporary system backup folder on the management server. You can select another temporary folder by right-clicking the notification area's Management Server service icon and by selecting Select shared backup folder.

System backup password

You can choose to protect the overall system configuration by assigning a system configuration password. After you assign a system configuration password, backups are protected by this password. The password settings are stored on the computer that is running the management server in a secure folder. You will need this password to:

- Restore the configuration from a configuration backup that was created with password settings different than the current password settings
- Moving or installing the management server on another computer due to a hardware failure (recovery)
- Configure an additional management server in a system with clustering

Scheduled backups

For large systems, Milestone recommends that you define scheduled backups. The management server stores your system's configuration in an SQL database. Milestone recommends that you regularly make scheduled backups of this SQL database as a disaster recovery measure. While it is rare to lose your system configuration, it can happen under unfortunate circumstances. Backups also have the added benefit that they flush your SQL database's transaction log.

Microsoft® SQL Server Management Studio, a tool downloadable for free from their website (<https://www.microsoft.com/downloads/>), is required for performing scheduled backup and restore. In addition to managing SQL Servers and their databases, the tool includes some easy-to-use backup and restoration features. Download and install the tool on your management server.

Manual backups

If you have a smaller setup and do not need scheduled backups, you can back up your system configuration manually. The following steps describe how to back up a system configuration manually:

1. From the menu bar, select **File > Backup Configuration**.
2. Read the note in the dialog box and click **Backup**.
3. Enter a file name for the .cnf file.

4. Enter a folder destination and click **Save**.
5. Wait until the backup is finished and click **Close**.

Monitoring the solution

The Computer Vision platform can easily have more than 50 VMs, each performing different roles. It is important to be able to monitor the core infrastructure and also the VMS or CV applications.

Topics:

- [VxRail platform monitoring](#)
- [BriefCam](#)
- [Ipsotek](#)
- [Milestone Systems](#)

VxRail platform monitoring

The VxRail platform has extensive monitoring options out of the box with vCenter. Any shortage of resources will be flagged as an alert and action can be taken.

To gain additional insights in to the infrastructure during our testing we used [vRealize Operations](#).

By using vRealize Operations we were able to easily get reports across multiple VMs and total system wide performance. The tool also supports analyzing GPU usage which is important in a multi GPU environment.

BriefCam

Large-scale deployments handle thousands of video hours per day and usually include 10 servers or more. The architecture may include infrastructure components specifically for monitoring, centralized logging, queuing, and more.

In all implementations, the VSServer service is responsible for various maintenance and monitoring related activities including:

- Watchdog the RESPOND tasks in case of a task failure
- Create RESPOND tasks when a user creates or modifies a rule
- Provide a live image for the RESPOND task configuration wizard
- Provide the list of cameras for the camera activation dialog of the Web Admin.
- Create the scheduled RESEARCH tasks
- Send the outbound alerts to the outbound API and also sends alerts to the VMS clients that have real-time alerts integration (level 2a or above)
- Trigger the data maintenance activity
- Clear inactive sessions

Ipsotek

The Ipsotek platform runs as a distributed system. The critical components for processing camera traffic are the Processing nodes. The Ipsotek Management node load balances workload among the available Processing nodes. In a Large deployment there could be 40 processing nodes. It is important to monitor that all Processing nodes are correctly joined to the cluster. Since the Processing nodes also contain GPUs, it is useful to monitor GPU utilization across the Ipsotek cluster.

In order to easily monitor the Processing nodes, the Ipsotek VConfigure administrator console has the option to get **Processing Node Information** in the **Server Settings** menu. This view lists the connected Processing nodes.

The Processing node information also covers the GPU utilization for each node that is processing camera traffic. The key metrics are:

- Hostname - The name of the processing node

- GPU % - The overall GPU percent utilization
- Free RAM MB - The amount of memory free on the GPU
- Decoder % - The utilization of the GPU hardware decoder
- Encoder % - The utilization of the GPU hardware encoder

In addition to the high-level view of the Processing nodes, a low-level monitoring portal is included in the Ipsotek deployment. This portal is based on Kibana and streams active events from all servers in the Ipsotek system. Any errors or warnings in any part of the system show up on the log monitoring dashboard. This is useful to trace a system issue or to monitor if a particular GPU is out of resources.

Milestone Systems

You can access your XProtect system from other computers using the clients. The XProtect Smart Client is used for viewing live videos on the Live tab, recorded videos on the Playback tab, and monitoring system status on the System Monitor tab. The Management Client is used for configuring and managing the system on other computers.

XProtect Management Client

Logs can be accessed from the Management Client by going to the Site Navigation pane and selecting Server Logs. Logs are used to get a detailed record of user activity, events, actions, and errors in the system.

Log type	What is logged?
System logs	System-related information
Audit logs	User activity
Rule-triggered logs	Rules in which users have specified the Make new <log entry> action.

- In each log window, you can apply filters to see log entries from, for example, a specific time span, a device, or a user.
- Exporting logs helps you to, for example, save log entries beyond the log retention period. You can export logs as comma-separated values (.csv) files.

XProtect Smart Client

The XProtect Smart Client has the following tabs:

Tab name	Description
Live	View live video
Playback	View recorded video
Search	Create advanced searches for video and metadata.
Alarm Manager	Investigate incidents and alarms
System Monitor	View system information

- The System Monitor tab displays a visual overview of the current state of your system servers, cameras, other devices, and the computer running XProtect Smart Client.
- For more information, see [System Monitor tab \(explained\)](#).

Troubleshooting recommendations

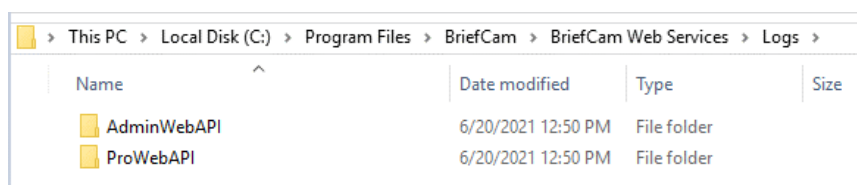
Topics:

- [BriefCam](#)
- [Ipsotek](#)
- [Milestone Systems](#)

BriefCam

Every service that is installed on a server creates a separate folder for writing log files. If more than one instance of a module is installed, a separate log folder is created which will have an increasing number value appended to its name. No folders or files are created for any module that is not installed. Log files are started fresh each day and organized by including their creation date in the filename when the daily log is closed. Log files do not store any sensitive data, such as passwords or access tokens. Administrators can also set a maximum log file size in case daily logs are getting too large to manage efficiently. There is also a setting to control when to trigger an automatic .zip archive creation based on the number of files created in a day.

Information useful for troubleshooting BriefCam's Web Services can also be found in the Windows Internet Information Services (IIS) logs. The location of these logs for a default installation can be found in the following path and folders.



Name	Date modified	Type	Size
AdminWebAPI	6/20/2021 12:50 PM	File folder	
ProWebAPI	6/20/2021 12:50 PM	File folder	

Figure 25. Windows IIS log location example

In addition to the BriefCam RESEARCH Service log files that are normally located at: `C:\Program Files\BriefCam\BriefCam Server\logs\BIRuleEngineService-O`, there are log files that are produced by the integration with the Qlik service. Qlik's trace log folder has separate log files for each element such as Engine and App Migration. See the *BriefCam Administrators Guide* for more detailed descriptions of troubleshooting tips using log files.

Ipsotek

A common scenario during our testing was to exhaust the system so that no additional camera workload can be processed. In this case, the status of the camera in the VConfigure tool is **Active - No Video**. In this state, the video stream is not being processed and events are not being reported. This status can also be visible if a camera stream is incompatible or was lost for some reason. To check if this is a temporal issue, it is possible to disable and then enable the camera in the VConfigure tool. If the status does not go to **Active - Streaming**, then go to Kibana and check the logs.

The Kibana dashboards show all log data from VMs that are part of the Ipsotek system. The following image is an example of a Kibana dashboard. The view is configurable to show any data that is part of the Ipsotek logs at any log level. The Ipsotek Support team should be consulted to set up the required log dashboards based on specific customers needs.

The screenshot shows the Elastic Logs interface. The top navigation bar includes 'Stream', 'Anomalies', 'Categories', and 'Settings'. Below the navigation bar, there's a search bar and filters for 'Customize', 'Highlights', and 'Last 1 day'. The main content area displays a table of log entries with columns for 'event.dataset', 'Message', 'source.application.name', and 'source.host.name'. The log entries show various system messages related to video processing and stream management.

event.dataset	Message	source.application.name	source.host.name
Apr 8, 2022	video_stream "subsession" was set up successfully (client ports were opened)	vps	ipso-ps-3
87:34:40.497	data sink for the "video/H264" subsession was created successfully	vps	ipso-ps-3
87:34:40.554	The started stream is too demanding for the remaining processing power of the GPU. Rejecting the stream...	vps	ipso-ps-3
87:34:40.557	Video Processing Thread finished	vps	ipso-ps-3
87:34:40.823	Pipeline 21 is not operating optimally. Restarting pipeline...	vps	ipso-ps-3
87:34:40.838	Finalising ...	vps	ipso-ps-3
87:34:40.839	Streamed closed	vps	ipso-ps-3
87:34:40.869	synchronizing and deleting streams...	vps	ipso-ps-3

Figure 26. Kabana dashboard example for Ipsotek VMs

Milestone Systems

Milestone provides online articles for guidance on troubleshooting the components of the XProtect VMS and the associated clients. The following list shows a few topics that are most relevant to this Design Guide:

- [Milestone Documentation / XProtect VMS products / XProtect VMS administrator manual / Installation log files and troubleshooting](#)
- [Milestone Documentation / XProtect VMS products / XProtect VMS administrator manual / Debug logs \(explained\)](#)

More troubleshooting topics can be found by using the [search tool](#) on the Milestone XProtect VMS home page.

Summary and conclusions

We found that the VxRail with VMware provided a reliable and scalable solution for hosting VMS and CV applications from different vendors on a common converged platform.

Dell Technologies, together with three market-leading VMS and CV vendors, have designed an integrated solution running on a common platform that is easy to install, easy to operate, and easy to grow.

The results of our HA testing show that virtualized instances of Milestone Systems XProtect for video management together with CV applications from both BriefCam and Ipsotek provide native high availability (HA) features and scale-out performance that are critical in an enterprise-class solution. The VxRail HCI system with VMware virtualization plus NVIDIA A40 server-class GPUs with vGPU provides additional features for managing high availability and scale-out needs that can complement the CV and VMS application's native features.

All of these products were installed on a single VxRail 5-node cluster for full system integration testing. Storage for the video management archive was provided by Dell PowerScale scale-out NAS storage. All video streams used for testing were provided by camera simulation software reading from prerecorded local video files in a continuous loop. We tested a maximum of 840 simulated camera streams being archived to the VMS together with each CV application performing real-time analytics on 320 camera streams. We also found that a single Ipsotek VM with one NVIDIA A40 GPU could process real-time analytics from 30 camera streams reliably. We also found that a single BriefCam VM with one NVIDIA A40 GPU could process real-time analytics from 50 camera streams reliably.

Finally, we identified an approach to help to size a platform of any size. The solution scales horizontally as more cameras are added to the system. We also defined rules around the placement of all VMs in the platform to ensure high availability is retained at all times.

References

The following resources provide additional relevant information:

Dell Technologies

- [Virtualizing GPUs for AI with VMware and NVIDIA - Design Guide](#)
- [VxRail home page](#)
- [Configuration Best Practices-Dell EMC Storage with Milestone XProtect Corporate](#)
- [Sizing Guide-Dell EMC Storage with Milestone XProtect Corporate](#)

NVIDIA

- [NVIDIA A40](#)

Milestone Systems

- [Milestone Systems home page](#)

Ipsotek

- [Ipsotek home page](#)

BriefCam

- [BriefCam home page](#)