

Dell EMC PowerEdge R740xd as a Dedicated Milestone Server, Using Nvidia GPU Hardware Acceleration

Dell IP Video Platform Design and Calibration Lab

H17250.1

Abstract

This document provides an overview of Nvidia GPU hardware acceleration using the Dell EMC PowerEdge R740xd as a dedicated Milestone VMS server. The tests include various scalable configurations across disk types, populated CPUs, the introduction of GPU, and how GPU scales.

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.



Topics:

- [Document overview](#)
- [Hardware overview](#)
- [Test setup](#)
- [Test results](#)
- [Findings and design considerations](#)

Document overview

This document provides an overview of the Dell EMC PowerEdge R740xd used as a dedicated Milestone video management software (VMS) server, with NVIDIA GPU hardware acceleration.

The R740xd is the latest generation Dell high disk intensity platform, which also provides support for GPU adapters. It is an ideal platform for the safety and security industry. Milestone has also introduced support for GPU decoding on their video management software. The tests include various scalable configurations across disk types, populated CPUs, the introduction of GPU, and how it scales.

This document focuses on testing with the following configurations:

- Fully populated SATA drive configuration, single CPU (No GPU)
- Fully populated SATA drive configuration, dual CPU (No GPU)
- Fully populated SAS drive configuration, dual CPU (No GPU)
- Mixed DISK, SSD and SAS, single GPU
- Mixed DISK, SSD and SAS, dual GPU

Dell IP Video Platform Design and Calibration Lab

In collaboration with Milestone Systems, the Dell IP Video Platform Design and Calibration Lab is an environment that provides a scalable controlled loading facility tailored to IP camera architectures. Platform design concepts can be implemented and calibrated to validate their potential capability.

Based at the Dell Customer Solution Centre in Ireland, using a network of global Customer Solution Centers, Dell helps customers strategize, design, validate, and build solutions.

Milestone Systems

Milestone Systems is the world's leading provider of open platform IP safety and security software. Milestone has provided easy-to-use, powerful video management software in more than 100,000 installations worldwide.

Milestone XProtect products are designed with open architecture and are compatible with more IP cameras, encoders, and digital video recorders than any other manufacturer. Because Milestone provides an open platform, you can integrate today's best business solutions and expand what is possible with future innovations. Go to www.milestonesys.com for more information.

XProtect Corporate

XProtect Corporate is IP-based video management software designed for large-scale and high-security installations. It is built with innovative technology designed to ensure end-to-end protection of video integrity and boost the overall performance of your system with hardware-accelerated video decoding.

In addition to the central management of all servers, cameras and users in a multi-site set-up, XProtect Corporate includes an integrated video wall for operators demanding situational awareness of any event. The software supports failover recording servers for mission-critical installations that require continued access to live and uninterrupted video recordings. Running on a recording engine with a recording rate of minimum 3.1 Gb/s, XProtect Corporate fully supports installations with 24/7 operation requirements, such as airports and casinos.

Release 2018 R1 implements support for NVIDIA GPU decoding.

Hardware overview

The R740xd server is an ideal platform to use as a dedicated video management system for Milestone VMS and provides the following benefits as a platform:

- The 2U platform is expandable up to 180 TB of local storage (18 x 3.5" SATA or SAS disks)
- Option to place the operating system on the Boot Optimized Server Storage (BOSS) PCIe adapter, in RAID 1
- OEM Ready

Additional features:

- Drive bays—Capable of hosting up to 18 x 3.5" high density disks, in either SATA or SAS 7.2K formats
- Storage controllers—Supports the Dell Boot Optimized Storage Solution (BOSS) PCIe adapter.
 - The Dell H740P PERC RAID controller has the performance benefit of 8 GB of non-volatile cache memory
- Processor—Dual Xeon scalable processors CPU support
- GPU options—Scalable GPU offerings
- Power supplies— Redundant power supplies
- Embedded management—Remote management leveraging the iDRAC, and Dell EMC OpenManage
- Security—Improved security, from a secure component supply chain, cryptographically signed firmware packages, and secure boot, Server Lockdown, and Secure Erase

OpenManage systems management

The Dell EMC OpenManage systems management portfolio includes powerful hardware and software management tools and consoles. OpenManage simplifies the lifecycle of deploying, updating, monitoring, and maintaining Dell EMC PowerEdge servers.

iDRAC Lifecycle Controller

The PowerEdge server provides storage expandability that allows you to adapt to your workload and operational demands. With comprehensive storage options, the server offers various drive types, internal and external storage controllers, and different backplanes for varied number of drives. The microcontroller is responsible for acting as an interface or gateway between the host system, such as server management software, and the peripheral devices. These peripheral devices, which may or may not be Intelligent Platform Management Interface (IPMI) compliant, consist of the power supplies, the storage backplane, integrated storage controllers, the control panel with semi-intelligent display, and Lifecycle Controller.

Our testing used the NVIDIA P4 Tesla adapters, which have a long product lifecycle and are:

- Designed for enterprise-grade reliability, including 24x7 operation
- Fully tested and qualified on the Dell R740xd server platform
- Includes H.264 and H.265 decoding, supported by Milestone
- Scalable to allow multi-GPU, multi-monitor, and additional GPU decoding acceleration

Table 1. Target platform

System configuration	R740xd low SATA disks	R740xd high SAS disks	R740xd GPU/SSD and SAS disks
CPU model	Intel Silver 4114 2.2 GHz (dual)	Intel Silver 4114 2.2 GHz (dual)	Intel Silver 4114 2.2 GHz (dual)
RAM	32 GB (2x16)		
Disks	BOSS + 18x4 TB 7.2K SATA	BOSS + 18x10 TB 7.2K NLSAS	BOSS + 12x10 TB 7.2K NLSAS and 6x800 GB SSD 12G SAS 10DWPD
PowerEdge RAID controllers (PERC)	H740P	H740P	H740P
NIC	Intel X500 10G+1350 1G/ +X710 10G SPF +		
PSU	750W x 2	750W x 2	1100W x 2
Graphic card	No	No	NVIDIA Tesla P4
OS	Windows Server 2016 St.		
OS disks	BOSS 2xM2 240G R1		
Live DB Partition and formatting	18x4 TB SATA - R6 (58 TB usable)	18x10 TB SATA - R6 (140 TB usable)	6x800 GB SSD (10DWPD) - R5

Table 1. Target platform (continued)

System configuration	R740xd low SATA disks	R740xd high SAS disks	R740xd GPU/SSD and SAS disks
Live DB config	RAID 6 / 1 MB Stripe / 64K Allocation Unit		RAID 5 / 1M Stripe / 64 Allocation Unit
Archive DB partitioning	—	—	12x10 TB SAS - RAID 10 (50 TB usable)
Archive DB config	—	—	RAID 6 / 1 MB Stripe / 64K Allocation Unit
Milestone XProtect Corporate release	2018 R1 build 7115	2018 R1 build 7115	2018 R2 12.2a build 8439
BIOS/Drivers	OS update + system updates with BIOS		

Nvidia adapter specification

The NVIDIA Tesla P4 is an enterprise grade GPU adapter that is designed for 24x7 operation, scalable up to six physical adapters in the R740 platform.

Our configuration of the R740xd has two available PCIe slots due to the high drive density, the BOSS PCIe drive adapter, and the 10G SPF+ network adapter.

The NVIDIA Tesla P4 is powered by the NVIDIA Pascal architecture and is purpose-built to boost efficiency for scale-out servers running deep learning workloads, enabling smart responsive AI-based services.

Features

The Nvidia adapter includes the following features:

- Small form-factor, 50/75-Watt design fits any scale-out server
- INT8 operations, which reduce latency by a multiple of 15
- A hardware-decode engine capable of transcoding and inferencing 35 HD video streams in real time

Table 2. Specifications

GPU architecture	NVIDIA Pascal
Single-Precision Performance	5.5 TeraFLOPS*
Integer Operations (INT8)	22 TOPS* (Tera- Operations per Second)
GPU Memory	8 GB
Memory Bandwidth	192 GB/s
System Interface	Low-Profile PCI Express Form Factor
Max Power	50W/75W
Enhanced Programmability with Page Migration Engine	Yes
Server-Optimized for Data Center Deployment	Yes
Hardware-Accelerated Video Engine	1x Decode Engine, 2x Encode Engine

Test setup

The test environment comprises two separate networks, one to emulate camera feeds, and the other to emulate client loading. Both emulation environments can be calibrated to induce precise scaling to the focus test platform. They are controlled by Milestone test automation tools, which can increase the loading while monitoring the key resources of the test platform. In this case, the video streams are calibrated to provide a 4.4 Mbit/s throughput that comprises a H.264, 1080p, feed at 30 FPS.

To represent a realistic use case both writing (recording) and reading (playback) are performed on the disk with the aim of having 10 percent playback of the write throughput. Milestone Server-side video motion detection is enabled for all cameras.

When the optimum level is reached, the Dell LiveOptics tool gathers the system performance information over a 24 hour period. LiveOptics is a vendor, hardware, and platform agnostic standard for IT professionals to record and communicate their achieved benchmarks, workloads, or support concerns to others to accelerate decision time and reduce risk. Furthermore, Milestone Performance and Microsoft Performance (perfmon) Resource Monitor tools monitor performance and frame loss.

We used the NVIDIA_SMI.exe and Microsoft PerfMon to monitor the GPU performance.

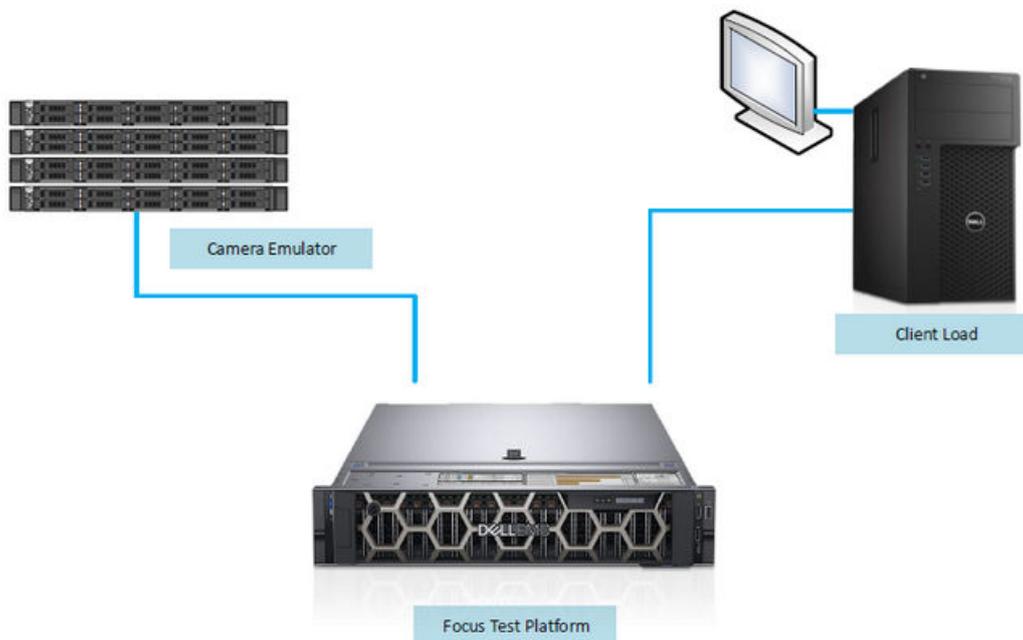


Figure 1. R740xd test configuration

Drive configurations

We installed the operating system and application to the BOSS adapter, using RAID 1, and M.2 SSDs for performance. As a dedicated VMS server, the IOPS to the operating system drive is high, which is why Dell EMC recommends fast disks for the OS/application drive.

- RAID** We implemented a RAID 6 on the database drive, to maximize the available space, while providing two points of failure on the RAID configuration.
- Virtual disk/volume** The PERC H740P virtual disk was configured with a stripe size of 1 MB. Formatting of the volume was with a 64K allocation unit.
- SSC** Using SSD disks as a stepping stone, running the Live Database on it, and archiving to the higher density disks provides the increased performance required during the archive process.
- Live Database size** Estimate the Live Database size to be ~ 50 percent of the total size of the SSD available space.
- Writes per day** Be aware of the number of writes per day. SSD disks have a performance metric of Drive Writes Per Day (DWPD). In the previous example, there is an available space of 4 TB, and an average of 11.28 TB daily writes, therefore mixed use 4DWPD disks are ideally suited to this application.

Test criteria

The following criteria were used in testing:

- Video images—1080p, 30 FPS, H.264 codec, 4.4 Mbit/s
- Video motion detect—On for all streams
- Read playback—Playback from recorded images at 20 percent of the write throughput
- System database—Disk space running at 80 percent full disk capacity, ensuring first in first out rotation
- The criteria for stopping the testing were:
 - Sustained media lost on input feeds
 - Sustained CPU performance is more than 70 percent
 - Sustained GPU utilization is more than 90 percent

Data points gathered

Data points were gathered for Microsoft Resource Monitor, application diagnostics, and GPU.

- Microsoft Resource Monitor**
- 24 hour LiveOptics analysis, at maximum performance
 - CPU percentage utilization
 - Memory usage
 - Network throughput on both the camera load interface and the client interface
 - Disk throughput
 - IOPS
 - Read/write ratio

- Application diagnostics measuring**
- Performance monitor, media lost in pipeline, across all cameras
 - System performance
 - Storage capacity usage

- GPU data points**
- Live monitoring, using the following command - `nvidia-smi dmon -i 0 -s u`
 - Perfmon—NVIDIA utilization

Test results

We followed best practices for normal situations that might arise in a standard production environment.

The server was loaded to represent a real world operating platform, running from a full database, using 80 percent of the total disk space. Motion detect was enabled across all streams, while also running a 20 percent client playback.

After calibration, we introduced multiple hard failures to the environment to monitor the impact on the performance of the application through the failure, and the recovery times to active recording and operation.

The camera emulator used is proprietary to Milestone, generating controlled feeds from a remote location to target platform. It ensures consistent, repeatable bandwidth in a controlled test environment.

Table 3. IP Video Lab results

		R740xd - SATA		R740xd SAS	R740xd	
		Single CPU	Dual CPU	Dual CPU	Single GPU	Dual GPU
	Camera emulator feed	StableFPS 1920x1080_Shopping_Mall_4Mbit_30FPS				
	Quantity	125	150	200	250	275
	Motion detection	Enabled	Enabled	Enabled	Enabled	Enabled
	Live DB size	48 TB	48 TB	117 TB	2 TB	2 TB
	Archive DB size	NA	NA	NA	40 TB	40 TB
	Database full (80% of total disc size)	Yes	Yes	Yes	Yes	Yes
	Observation period	24 hours	24 hours	12 hours	12 hours	12 hours

Table 3. IP Video Lab results (continued)

		R740xd - SATA		R740xd SAS	R740xd	
		Single CPU	Dual CPU	Dual CPU	Single GPU	Dual GPU
	Database	Live only	Live only	Live only	Live and Archive	Live and Archive
Platform						
	CPU load (Peak/Min)	77%	47%	54%	28%	25% / 28%
	Memory load sustained	9.98 GB	13.3 GB	15.1 GB	20 GB	21 GB
	Peak memory	12 GB	20.43 GB	19.21 GB	29.77 GB	21.41 GB
	GPU Decoder load (ave)	NA	NA	NA	90% +	70% + & 70% +
	GPU Decoder memory load	NA	NA	NA	~10%	~10%
Performance						
	Network BW (video-received)	522 Mb/s	637.18 Mb/s	841 Mb/s	1.1 Gb/s	1.3 Gb/s
	Network BW (playback-sent)	105 Mb/s	103.44 Mb/s	55 Mb/s	100 Mb/s	100 Mb/s
OS disks						
	OS drive IOPS @95%	3139	3690	5270	2500	833
Live DB disks						
	Live DB drive read/write ratio (IOPS)	54% / 46%	42% / 58%	55% / 45%	90% / 10%	92% / 9%
	Live DB size	48.95 TB (full) 84%	48 TB (full)	117 TB (full)	1.96 TB (full) 54% disk	1.96 TB (full) 54% disk
	Live DB drive disk throughput	80 MB/s	95 MB/s	113 MB/s	278 MB/s	354 MB/s
	Live DB average daily writes	5.9 TB	6.07 TB	8.01 TB	10.4 TB	11.28 TB
	Live DB drive IOPS @95%	131	145	107	1463	2675
Archive DB disks						
	Archive R/W ratio	NA	NA	NA	20% / 80%	23% / 77%
	Archive DB size	NA	NA	NA	39.16 TB (full DB) 75% disk	38.13 TB (full DB) 71.3% disk
	Archive DB drive disk throughput	NA	NA	NA	150 MB/s	223.8 MB/s
	Archive DB average daily writes	NA	NA	NA	10.4 TB	11.28 TB
	Archive DB drive IOPS @95%	NA	NA	NA	120	210
Disruptive application						
	Drive rebuild affect (RAID)	Negligible	Negligible	Negligible	Negligible	Negligible
	Live DB recovery time from power down	NA	<14 min	<20 min	<1 min	<1 min
	Archive DB recovery time from power down	NA	NA	NA	<12 min	<12 min
	Frames lost	~0	~0	~0	~0	~0

Comparative test results

The comparative test results are shown in the following graphics.

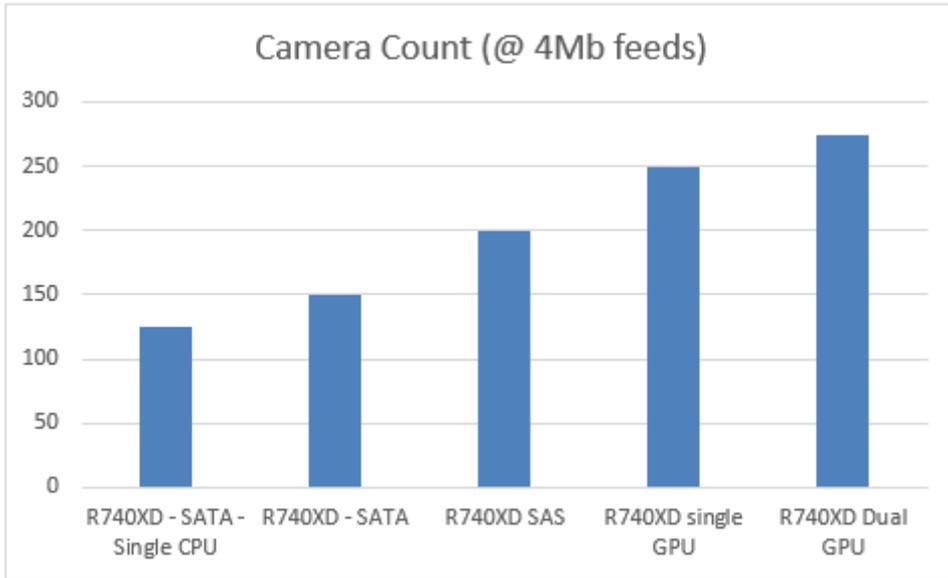


Figure 2. Camera count

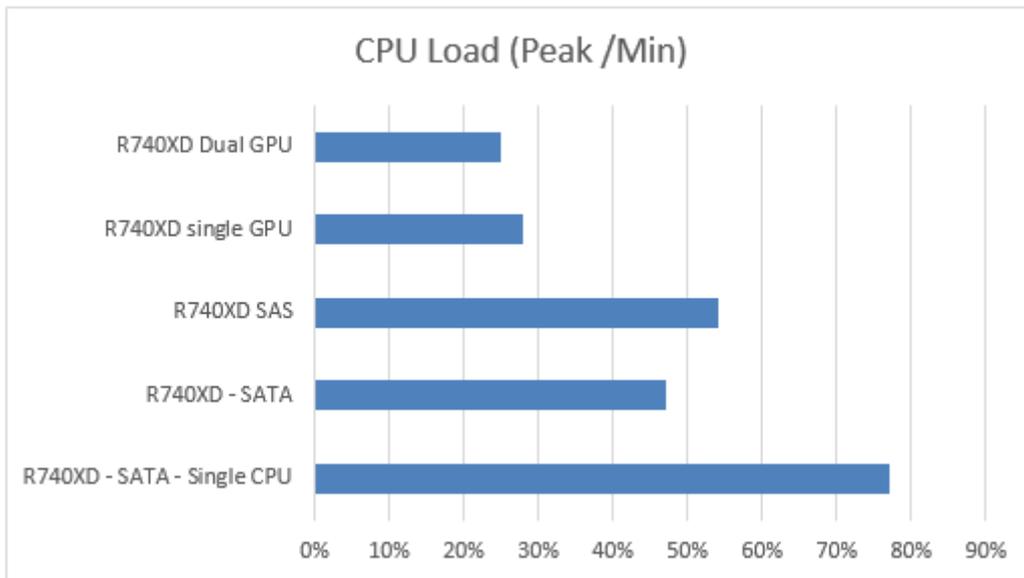


Figure 3. CPU load

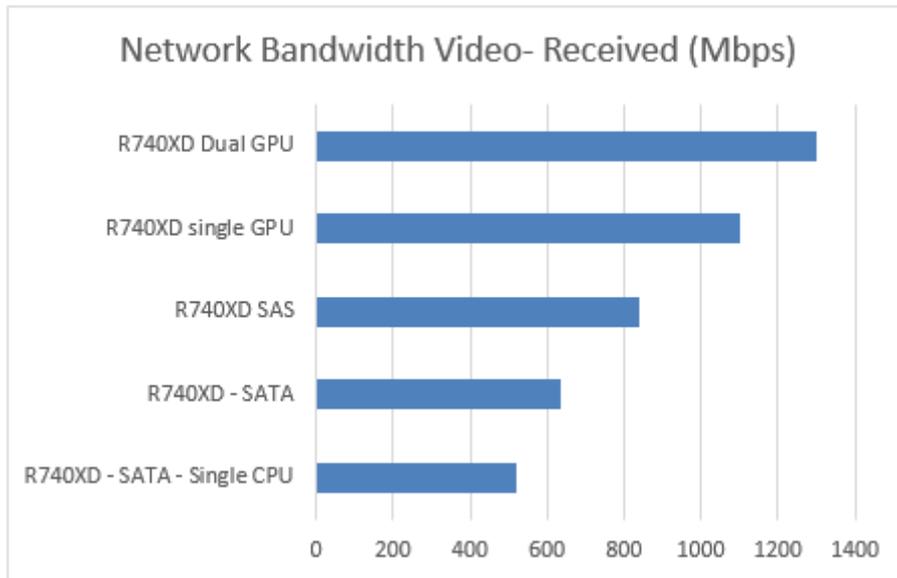


Figure 4. Network bandwidth

R740xd SATA single CPU

The R740xd SATA single CPU test results are shown in the following graphics.

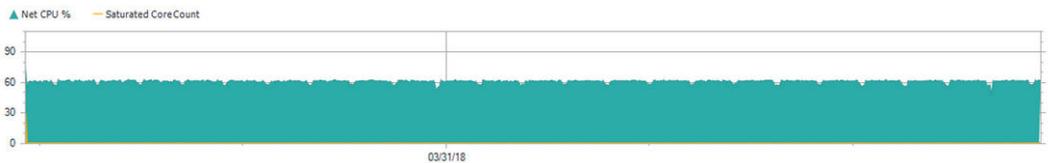


Figure 5. CPU performance

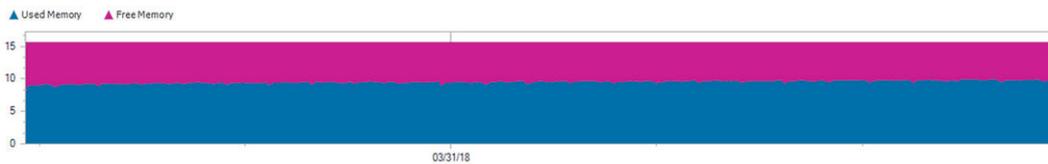


Figure 6. Memory performance

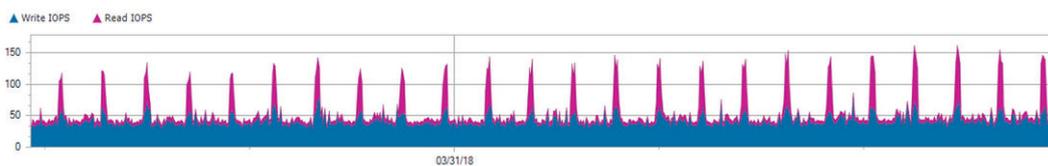
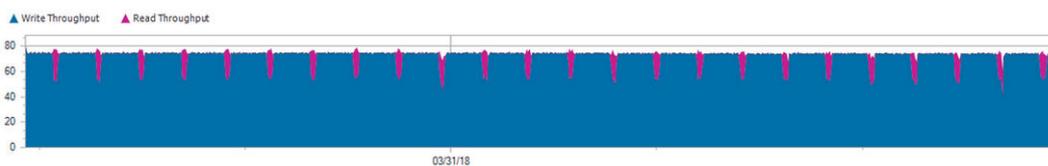


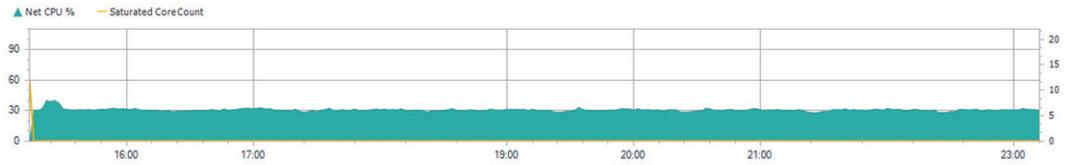
Figure 7. Live database drive performance



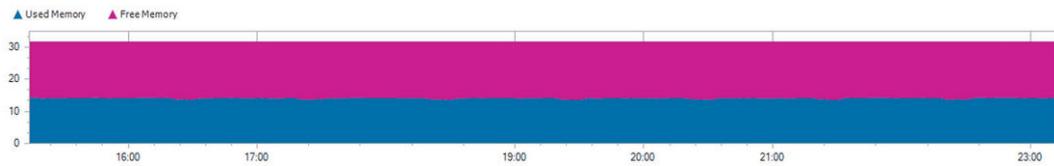
R740xd SATA dual CPU

The R740xd SATA dual CPU test results are shown in the following graphics.

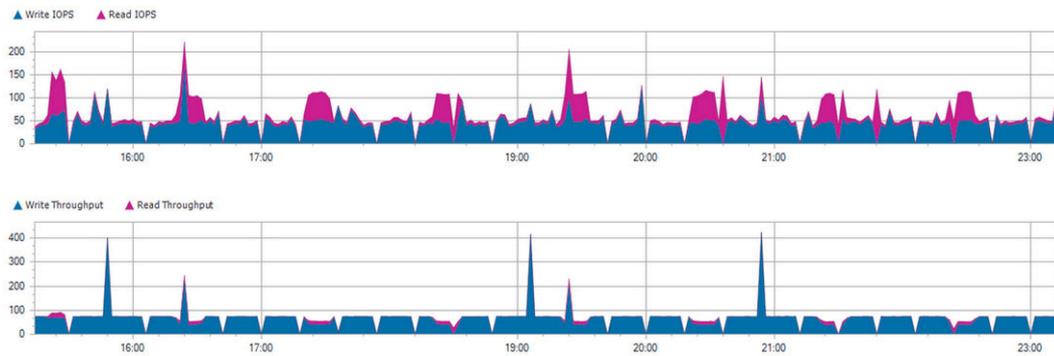
CPU performance



Memory performance



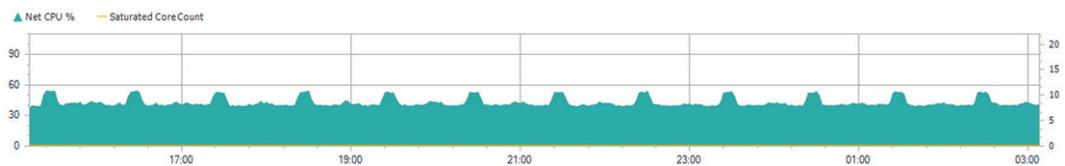
Live database drive performance



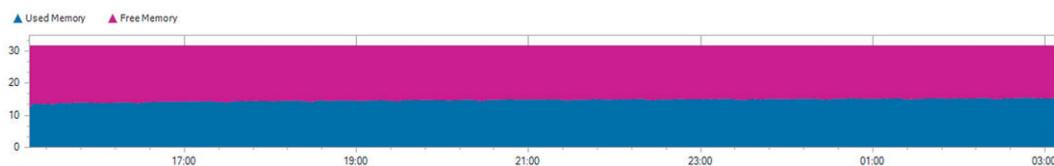
R740xd SAS dual CPU

The R740xd SAS dual CPU test results are shown in the following graphics.

CPU performance



Memory performance



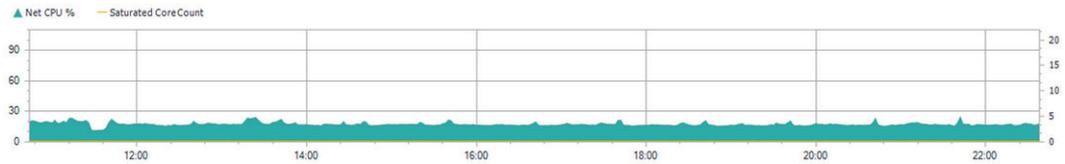
Live database drive performance



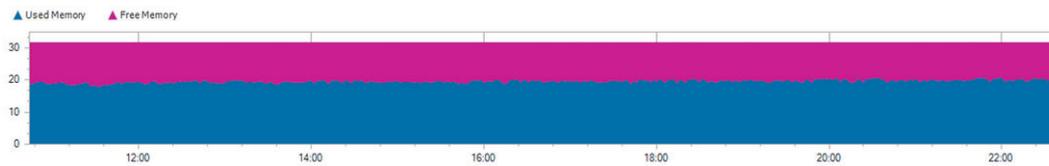
R740xd single GPU

The R740xd single GPU test results are shown in the following graphics.

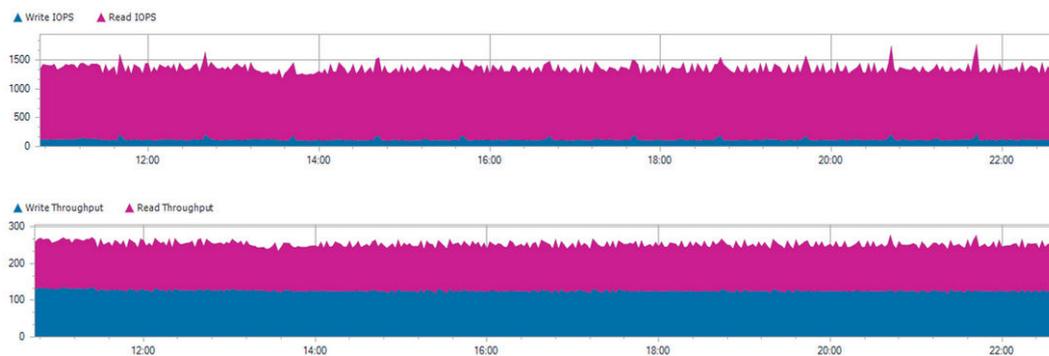
CPU performance



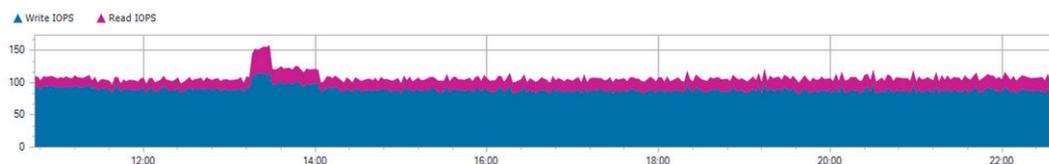
Memory performance

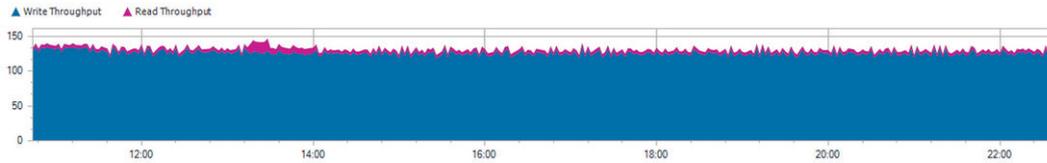


Live database drive performance



Archive database drive performance

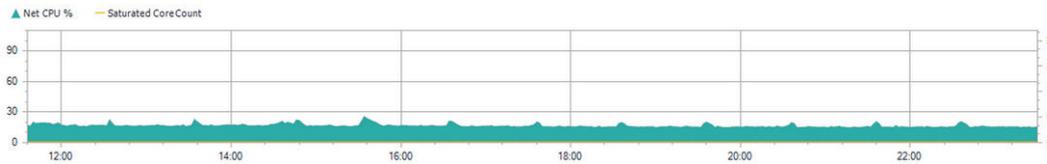




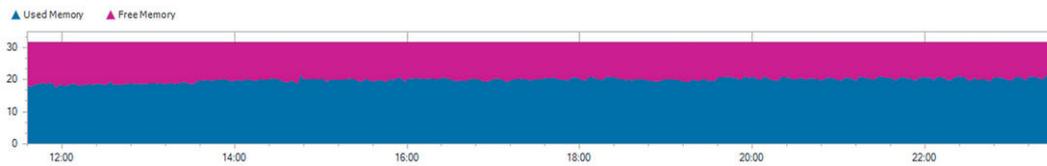
R740xd dual GPU

The R740xd dual GPU test results are shown in the following graphics.

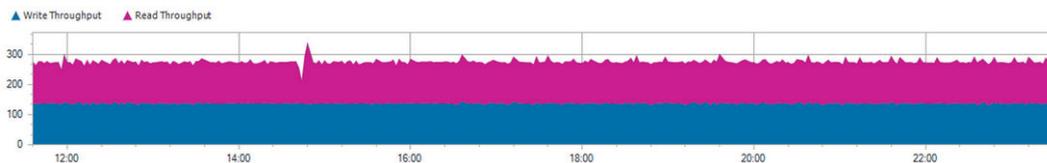
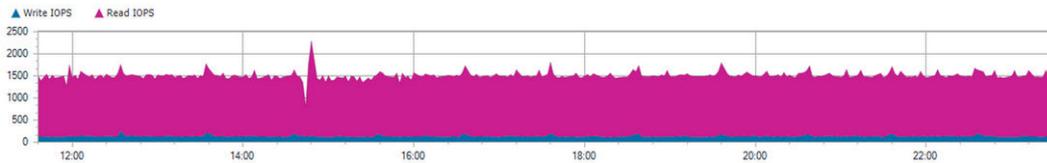
CPU performance



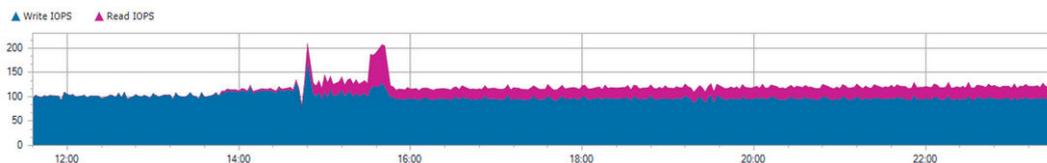
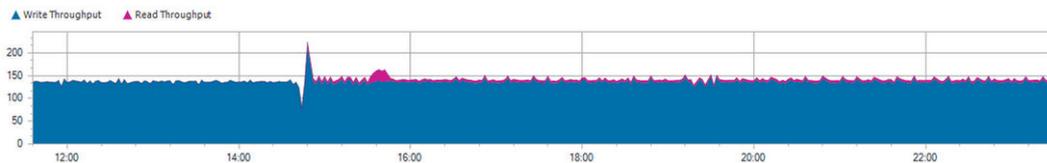
Memory performance



Live database drive performance



Archive database drive performance



Findings and design considerations

Performance

IOPS to the database drive plays an important part in determining the number of feeds, or the amount of bandwidth, to the server.

- To maximize the retention time space available, compared to redundancy, the 18 disks are configured to RAID 6.
- We achieved approximately a 33 percent increase in performance between the SATA 7.2K disks, and SAS 7.2K disks in the same configuration.
- If a higher performance is required on the database, RAID 10 can be implemented.
- Using SSD as a live volume step to archiving gives you the appropriate performance required during the archiving process.
- The Read IOPS quantity is roughly 10 times the Write IOPS size from the Live Database during an archive operation.
- Remember to size the archive volume with an appropriate speed to accept the higher IOPS. Testing with RAID 6 as a landing zone for archiving gave us no improvement on the incoming bandwidth. Changing to RAID 10 gave us the bandwidth required on the archive volume. SSD disks can also be considered if higher performance is required.
- On all tests, we used 10G networking for the video network.
- A key benefit observed during tests of using a relatively small Live Volume, on fast disks, is the recovery of the platform to recording from a hard power down. The recorder server took less than one minute to start recording again, while repairing the archive database in the background, compared to more than 20 minutes on the server with a single Live Database of 117 TB.

Disks

- SATA disks: Useful for the higher retention time.
- SAS 7.2K disks: Will increase the performance while retaining the same retention time.
- SSD: Ideal for a Live Database small volume, archiving to higher density disks.

GPU support

The R740XD natively supports the NVIDIA P4 GPU, scaling to a maximum of six adapters in the R740 model. The BIOS support of the adapter ensures appropriate power and cooling is provided to the adapter.

Milestone primarily uses the Decoder chip on the adapter. Also, the level of CUDA and memory leverage with the P4 adapters is approximately 10 percent with motion detection enabled. Analytics applications on the same server can use the extra bandwidth.

- Pascal-based adapters are required to support the decoding of H.264 and H.265 feeds.
- The load on the Decoder chip is balanced as you scale the quantity of adapters in the platform.
- Adding the GPU adapter increases the overall platform performance, in our case by 37 percent in the dual GPU configuration, while decreasing the CPU load by approximately 50 percent.

Milestone application recommendations

Dell EMC recommends the following configuration for the Milestone application:

- Use 200 cameras for each recording server instance (to a maximum of 300)
- Install a full SQL Server (standard or enterprise) edition when using more than 300 cameras in a system
- When creating various recording server instances, you must give each instance a valid IPv4 and IPv6 (IPv6 cannot be looped with ::1)