CoreStation MX

Dell EMC MX7000 incorporating Blades servers with Nvidia GPU hardware Acceleration in a chassis-based solution

Reference Architecture

Revision Number:0.1

Last update: 15th October 2021

Contents

1 About this document

1.1 Contacts

Name	Role	Details
Andrew holloway	Author	Andy.holloway@amulethotkey.com

1.2 Version control

Version	Date	Details
0.1	29/01/21	Final

2 Document Overview

This document is designed to provide an overview of the Dell EMC PowerEdge MX7000 Modular chassis, The Chassis consists of Dell EMC MX740C compute blades which are leveraging Nvidia T4 GPU's which are housed in the Amulethotkey Core Module Expansion System.

PowerEdge MX, a unified, high performance kinetic infrastructure, provides the agility, resiliency and efficiency to optimize a wide variety of traditional and new, emerging data center workloads and applications. With its kinetic architecture and agile management, the MX portfolio dynamically configures compute, storage and fabric, increases team effectiveness and accelerates operations. Its responsive design delivers the innovation and longevity customers of all sizes need for their IT and digital business transformations. PowerEdge MX ecosystem consists of a new chassis infrastructure, compute sleds, fabric switches, and a storage sled, all managed by Dell EMC. This make the MX7000 the ideal platform for any business wanting to run an efficient effective and scalable surveillance system, under one single chassis customer can run not only their Milestone VMS system but also incorporate any addition ISV platforms needed to enhance their requirements.

Testing

This document will focus on the following testing configurations

- Milestone Video Management system using a live Local Database.
- Milestone Video management system Archiving to an external storage system
- Milestone Video Management system direct write to and external storage array

The focus of the Testing was to confirm that the T4 GPU's in the servers were being leveraged by the Milestone software and to show what impact this has on the performance of the computer blades. A relatively low number of Camera load was used only 150 x 4Mbps feeds were used across the tests.

** It must be noted that the CPU and Memory in the Blades Far exceeded those needed for the workloads tested.

Using the Dell computer Vision Labs

For this testing Amulethotkey and Dell worked in collaboration with Milestone systems. The computer vision Lab is an environment that provides a scalable controlled loading facility that is tailored to IP camera architectures. Based in the Dell Customer Solutions centre in Ireland, helps customers and partners strategize, architect, validate and build out validated solutions.

3 Solution

3.1 Overview

3.1.1 Solution components

Objective

This Reference Architecture describes a solution to create a high-quality, performant and responsive desktop experience for two workload types - virtualized and dedicated trader workstation.

Technology

This solution utilizes several advanced technologies such as Dell EMC's MX modular architecture incorporating Amulet Hotkey's CoreModule, each of these components is optimized to deliver the bestpossible experience to end users and at the same time offer administration and support teams a means to maintain, configure and deploy easily and effectively.

The PowerEdge MX740c compute sled forms the foundation of the solution. Configure each compute sled with up to two NVIDIA T4 GPUs.

3.1.2 Software components

Milestone Systems

Milestone Systems is the world's leading provider of open platform IP video surveillance 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 supreme situational awareness of any event. The software supports failover recording servers, making it the perfect choice for mission-critical installations that require

continued access to live and uninterrupted video recordings. Running on the industry's best performing recording engine with a recording rate of minimum 3.1 Gb/s, Xprotect Corporate is ideal for installations with 24/7 operation requirements, such as airports and casinos.

3.1.3 System diagram

Figure 1 shows a high-level schematic of the deployment used.

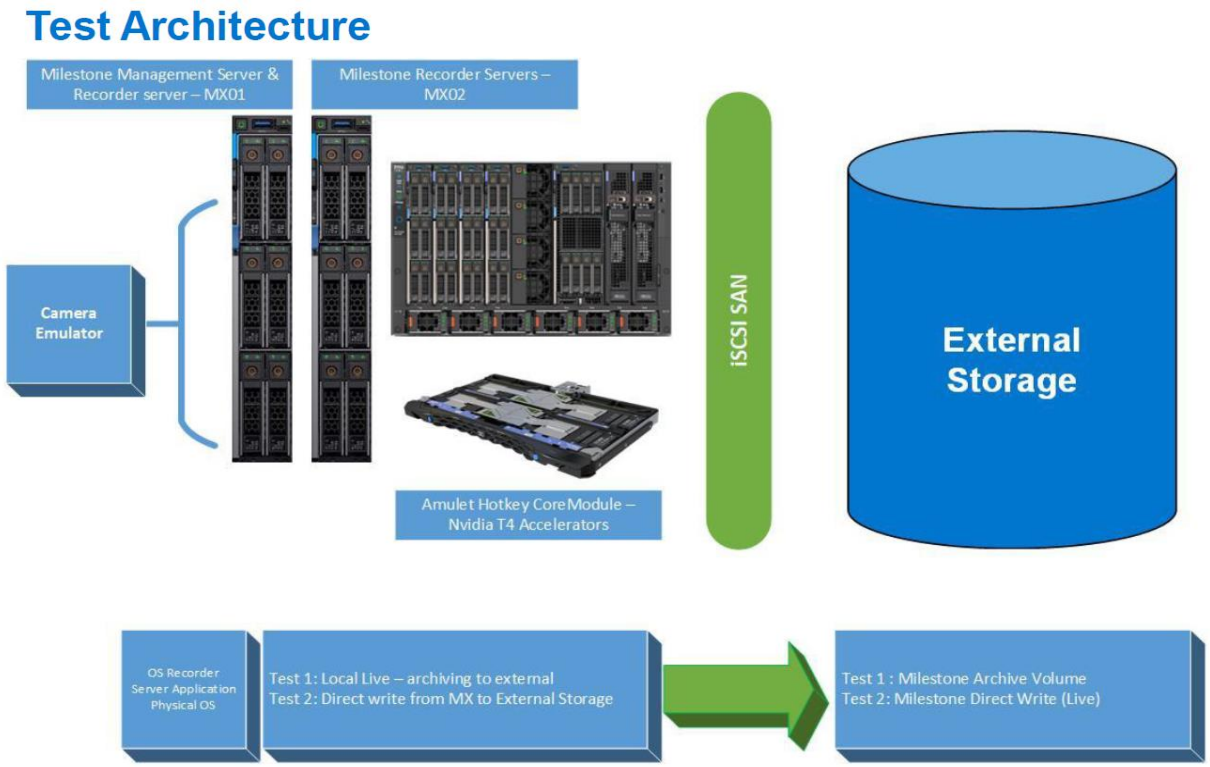

Figure 1: High-level solution example overview

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 Live Optics tool gathers the system

performance information over a 24-hour period. Live Optics 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.

3.2 Hardware solution components

3.2.1 PowerEdge MX7000 Modular Chassis

The PowerEdge MX7000 Modular Chassis is a 7U form factor modular chassis that accommodates various compute and storage sled combinations. These combinations are connected by high-speed connectivity fabrics, share power and cooling, and are managed by embedded systems management.

Figure 3 shows Fabric A networking using dual MX9116 FSE IOM leveraging Smart Fabric. This configuration provides an easy to manage single chassis deployment.

Figure 3: PowerEdge MX7000 Modular Chassis

MX7000 features

The PowerEdge MX7000 Modular Chassis has the following features:

- 7U modular enclosure with eight slots holds eight 2S (single), or four 4S (double-width) compute sleds;
- 25Gb Ethernet on sled, 12Gb SAS and 32Gb Fibre Channel I/O options;
- Single management point for compute, storage, GPU acceleration and networking;
- High-speed technology connections to uplink to enterprise networks;
- two fabrics for IO – Fabric A and Fabric B.

Chassis configuration

Each chassis is configured with Open Manage Enterprise – Modular Edition, which runs on the MX9002m management module. Up to 10 chassis (up to 80 MX740c sleds) can be managed from one 'lead' chassis, enabling you to deploy firmware/settings and manage alerts without having to visit each individual web interface.

CoreModules

For a CoreStation MX deployment, the Amulet Hotkey CoreModule GPU solution (see 3.3.4) installs in Fabric B, leaving Fabric A for breakout to the wider network.

An additional Fabric C provides storage IO capabilities.

Achieve additional scalability by using the MX9116 Fabric Switching Engine and MX7116 Fabric Switching Engine. Figure 4 shows Fabric B with dual CoreModules, with one installed and one removed.

Note: Any compute sleds that connect to CoreModule GPUs must be dual CPU.

Figure 4: MX7000 chassis with dual CoreModules

3.2.2 MX7000 chassis network

This solution validation has been performed with the following hardware configuration:

Top of rack: Dell EMC Networking S4128T (10 GbE ToR switch).

The S4128T switch optimizes the network for virtualization and remote desktop delivery with a high-density, ultra-low-latency ToR switch that features 24 x 10 GBase-T and 2 x QSFP28 ports.

MX7000 chassis: Dual MX5108n in single chassis configuration.

The MX5108n offers excellent performance and configuration options. Another recommended option for network scalability and management include the MX9116 and MX9117 switch modules.

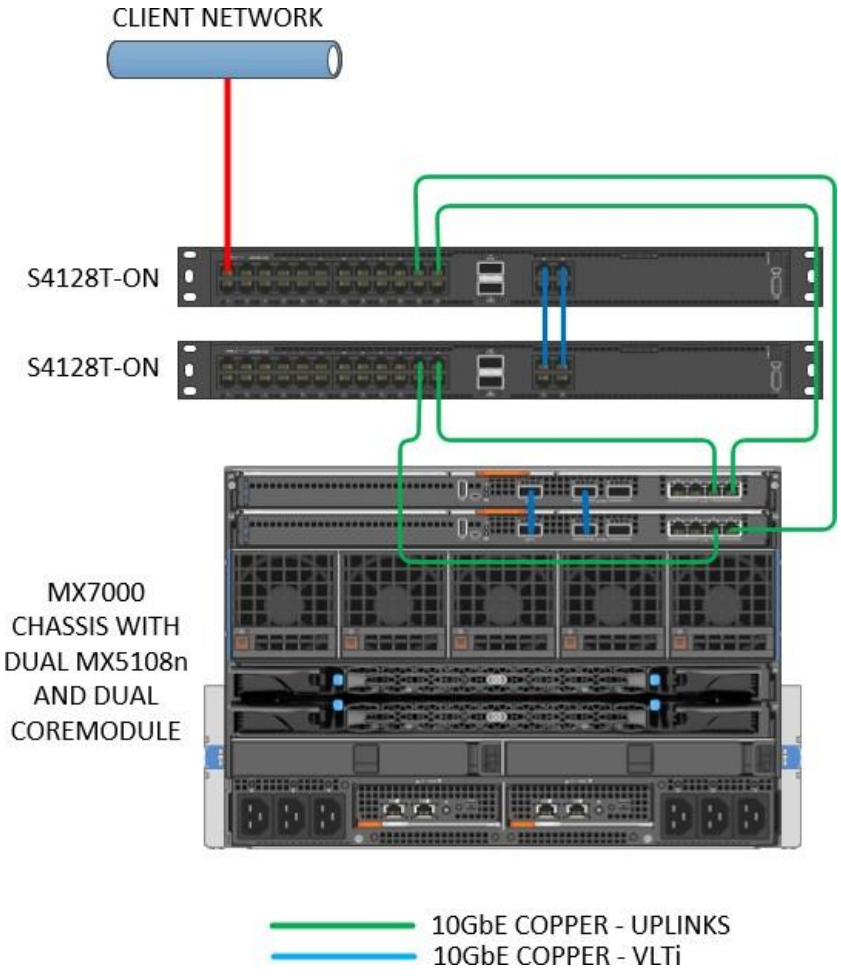

Figure 5 - Deployed configuration - network diagram

3.2.3 PowerEdge MX740c compute sled

MX740c overview

Each MX740c compute sled can support:

-up to two 28-core 2nd Generation Intel Xeon Scalable processors;

-24 DDR4 DIMM slots;

-up to six 2.5-inch SAS/SATA (HDD/SDD) drives;

-a maximum of eight MX740c servers in the PowerEdge MX7000 enclosure.

Each MX740c is equipped with iDRAC9, an advanced out of band management system.

The iDRAC controller has its own processor, memory, network connection, and access to the system bus. Key features include power management, virtual media access and remote console capabilities, all available through a supported web browser or command-line interface.

Amulet Hotkey mezzanine card (CoreCard)

A CoreStation MX740A (PowerEdge MX740c with an Amulet Hotkey Mezzanine card) within the PowerEdge MX7000 Chassis communicates via the mezzanine card (see Figure 6) to an Amulet Hotkey CoreCartridge in the CoreModule. The CoreCartridge houses two NVIDIA T4 GPUs. The Amulet Hotkey Mezzanine card enables a direct GPU PCIe pass-through connection to the compute sleds.

Figure 6: Amulet Hotkey Mezzanine card

3.2.4 Amulet Hotkey CoreModule

CoreModule overview

Developed by Amulet Hotkey, the CoreModule (see Figure 7) is designed exclusively for the PowerEdge MX7000 Chassis. This unique form factor enables the installation of multiple GPUs into the rear of the enclosure. The CoreModule is an essential component in enabling the MX7000 to support virtual workstations.

Each enclosure can support up to two CoreModules, each CoreModule can be configured with up eight NVIDIA T4 GPUs. A fully populated MX7000 can utilize up to sixteen GPUs – two per CoreStation MX740A. This density and balance of CPU to GPU is unique in the market and designed specifically with trading floor use cases in mind.

As well as enabling the MX7000 to support high-performance virtual workstations the CoreModule can also be used for several other GPU accelerated workloads:

- Virtual desktop and application acceleration;
- Artificial intelligence, machine learning and deep learning;
- Data analytics;
- HPC simulation and modelling;
- Rendering and video processing.

With no midplane between the CoreStation and the CoreModule, each GPU connects directly with PCIe 3.0 to the server, ensuring optimum performance and minimum latency.

Figure 7: Amulet Hotkey CoreModule

CoreModule configuration

The MX7000 accommodates two CoreModule expansion modules. These are installed in Fabric B only. Each CoreModule can host up to four CoreCartridges with a CoreCartridge consisting of two T4 GPUs. Each sled which requires a GPU must also be fitted with an Amulet Hotkey Mezzanine card.

Note: Compute sleds that do not require a connection to a GPU can occupy the same chassis as those that do.

3.2.5 NVIDIA T4 GPU

NVIDIA's T4 GPU (see Figure 8) accelerates a range of workloads, including high-performance computing, deep learning training and inference, machine learning, data analytics, and graphics. Based on Turing architecture, the T4 is optimized for mainstream computing environments and features multi-precision Turing Tensor Cores and RT Cores.

Figure 8: NVIDIA T4 in the CoreCartridge

NVIDIA T4 specifications

Turing Tensor Cores	320
CUDA Cores	2560
Ray Tracing RT Cores	40
Compute Performance	8.1 TFLPS Single-Precision, 65 TFLOPS Mixed-Precision, 130 TOPS INT8, 260 TOPS INT4
GPU Memory	16 GB GDDR6 (300GB/sec)
vGPU Profiles	1 GB, 2 GB, 4 GB, 8 GB, 16 GB
System Interface	x16 PCIe 3.0 (x8 PCIe 3.0 used for up to 16 GB/sec interconnect bandwidth)
Compute Platforms	CUDA, NVIDIA TensorRT™, HPC Applications and learning frameworks
Thermal Solution	Passive
Power	70 W

3.2.6 Hardware specifications

See below for key hardware components.

PowerEdge MX740c – Infrastructure Server

Component	Configuration
RAM	12 x 32GB DDR4 (384GB total)
CPU	2 x Intel(R) Xeon(R) Gold 6226 CPU @ 2.90GHz
Fabric A (Network)	Mellanox ConnectX-4 LX 25GbE KR Mezz Card

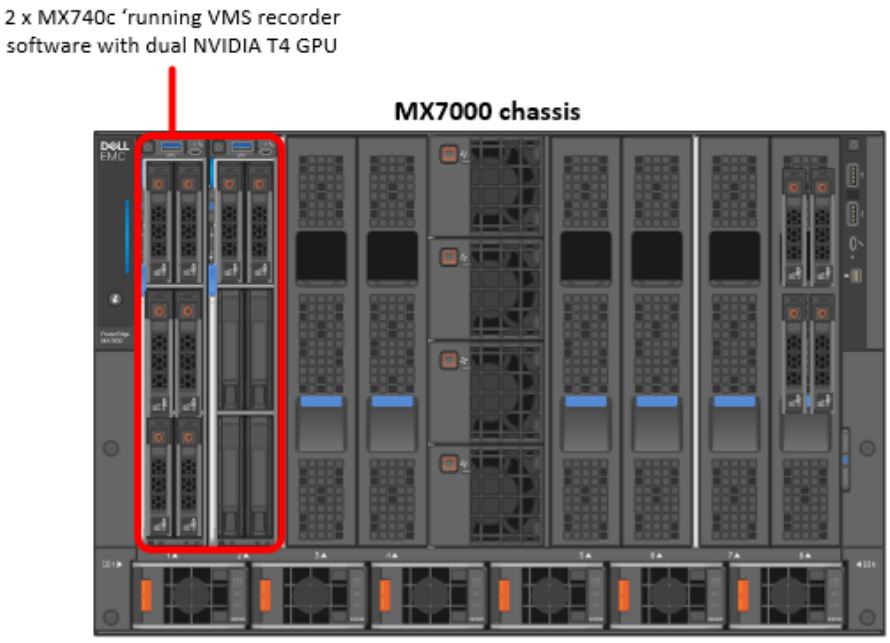

Figure 10: Reference chassis configuration indicating position and role of installed MX740c

Amulet Hotkey CoreModule

Component	Configuration	Enclosure Location
GPU	1 x CoreCartridge with dual NVIDIA T4 GPU	Fabric Slot B1
GPU	1 x CoreCartridge with dual NVIDIA T4 GPU	Fabric Slot B2

3.3 Software solution components

Component	Configuration
XProtect	Xprotect Corporate Version 2020 R3 Version 20.2a Build 4122
Server OS	Windows server 2019

3.4 Hardware

Below is a complete list of hardware used as part of this deployment. Each component is listed, along with a brief description of its function. All of the below are required as part of this solution deployment. For more detail on the specific versioning of these components, please refer to the Configuration Guidance documentation links in section 6.

Component	Function
Amulet Hotkey CoreModule	Hosts GPU accelerators for installation into the PowerEdge MX7000 Modular Chassis
Amulet Hotkey CoreCartridge	Installed within the CoreModule, each CoreCartridge hosts two NVIDIA T4 GPU
Amulet Hotkey Mezzanine Card	Installed into each compute node, in Fabric B, which requires a connection to a GPU
PowerEdge MX740c Compute Node	Dual CPU compute node, optimized to provide performant, scalable server functionality
PowerEdge MX7000 Modular Chassis	Contains all compute, IO, GPU and management components
Dell EMC MX5108n Ethernet Switch	Switch modules hosted within MX7000 chassis
Dell EMC Networking S4128T-ON	Top of rack switching
Dell 10G iSCSI /10TB	External Storage

4 Testing

4.1 Testing setup

The Test environment consists of 2 x MX740C servers, with a dedicated Nvidia T4 GPU adapter assigned to each server through the Amulet CoreModule. Server 1, MX01, hosted the full installation of the Milestone software. Server 2, MX02, was set up as a single Recorder Server.

We performed 2 tests.

Test 1: Taking the camera Emulated feeds into each server, test 1 stored the Live volume on a local SSD in the server, and then Archived to an external iSCSI storage repository.

Test 2: Took the same feeds, and directly wrote to a single Live Database on the external storage.

Each test was performed with the Nvidia T4 GPU enabled and disabled to take a comparison.

Focus for the results from the above is on the Acceleration differences and performance Gains

** Note the Specification of the CPU's and Memory in the test Servers are well above the required levels for Milestone.

The focus of the testing is to compare the Off Load capabilities of the T4 for the Decoding of the streams.

4.2 Tools used for Analysis

Software Tool	Function
LiveOptics Optical Prime	Used to gather all the perf data from the recorder server over 24 hrs
Microsoft Performance Monitor	To monitor the VideoOS recording severs Pipeline\Medias Lost/sec per camera
NVIDIA GPU Profiler	Monitoring GPU stats
Milestone Server Monitor	Monitoring Server Logs
Milestone Smart Client	To review camera footage and live monitoring

4.3 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

5 Test Results

5.1 Local Live Write, Archiving to External -MX01

GPU

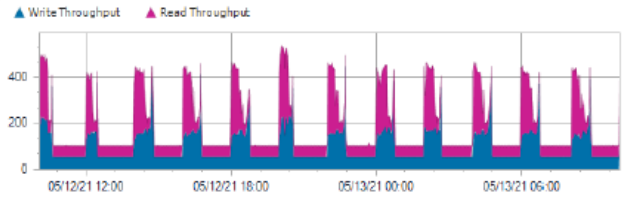

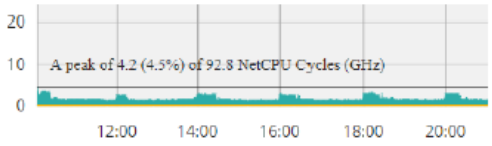

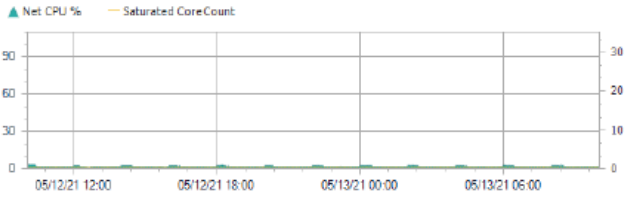

No GPU

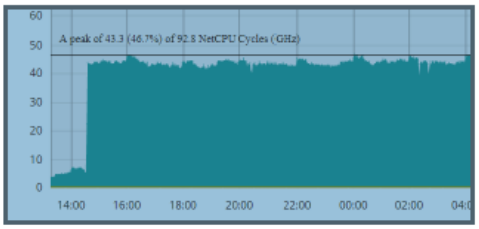

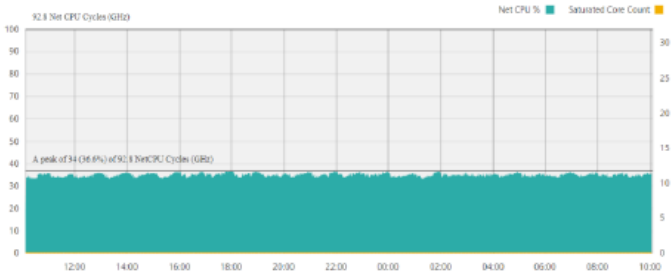

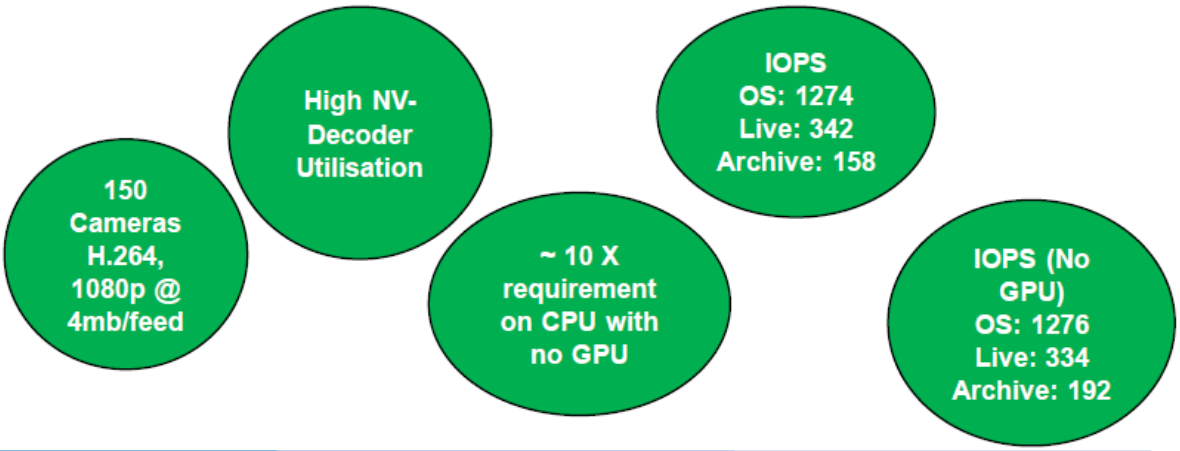

5.2 Local Live Write, Archiving to External -MX02

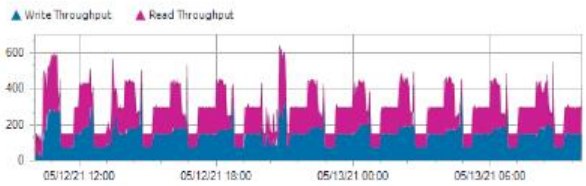

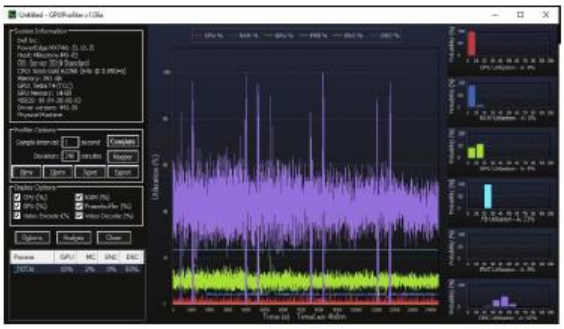

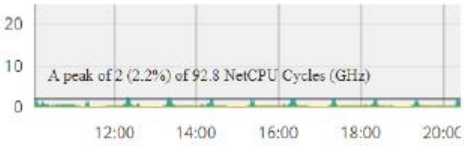

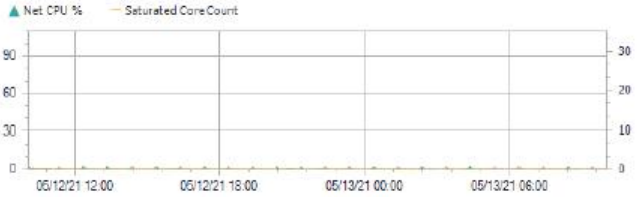

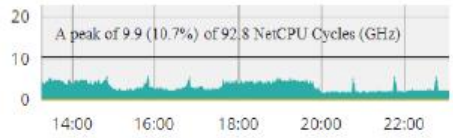

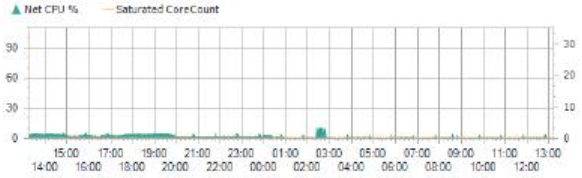

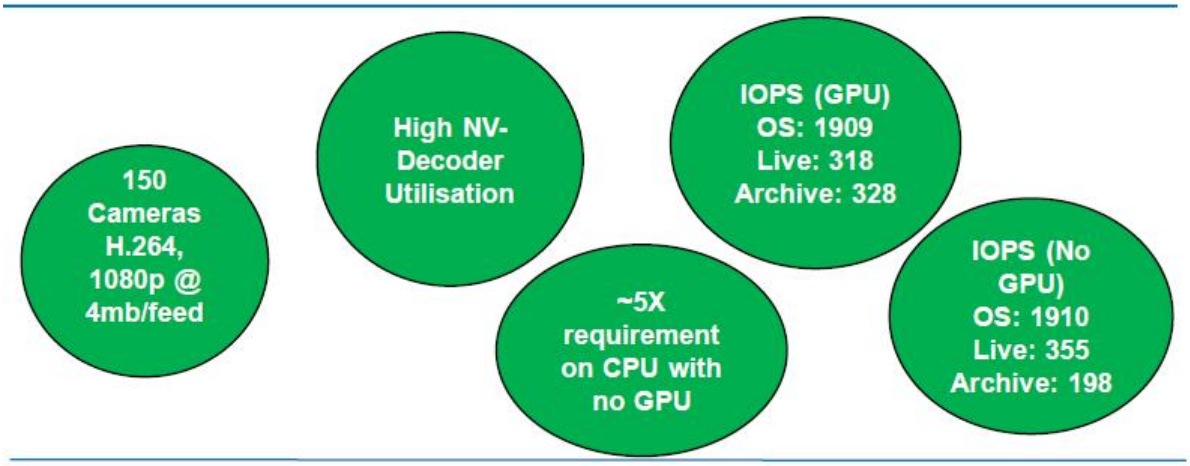

5.3 Direct Write to External Live Volume -MX01

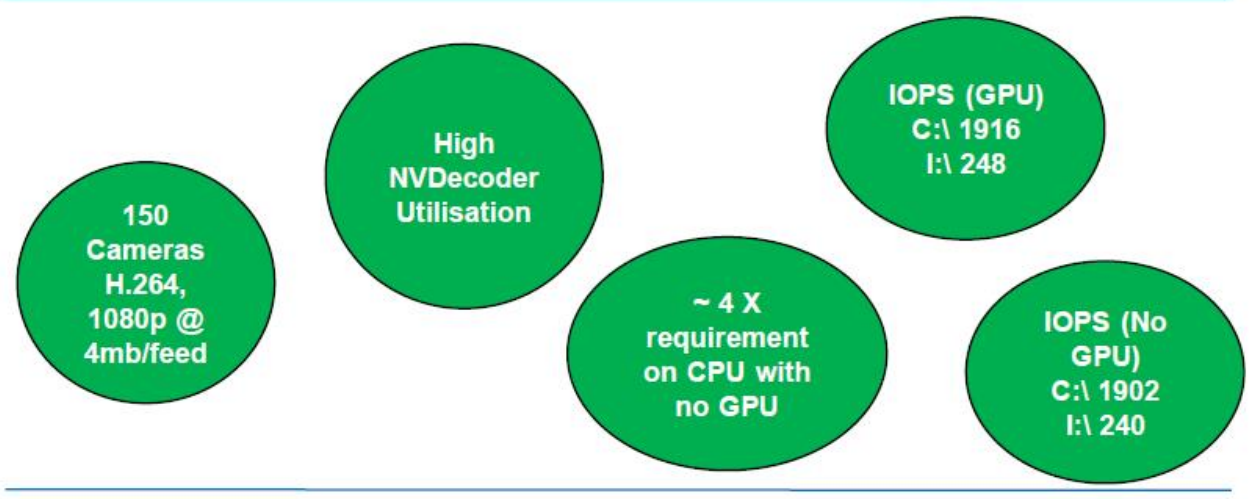

5.4 Direct Write to External Live Volume – MX02

GPU

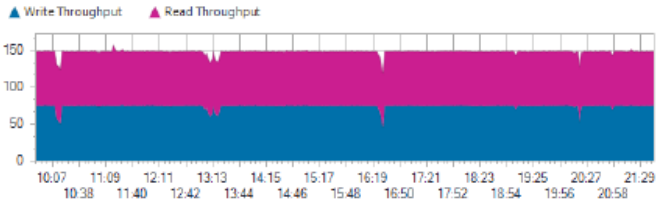

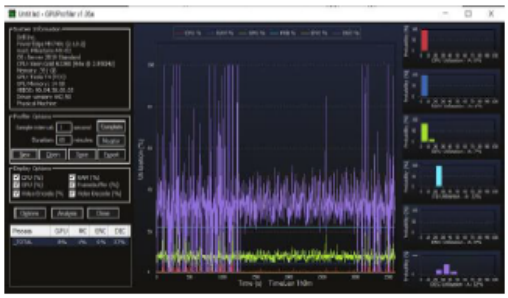

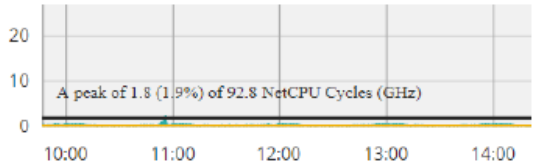

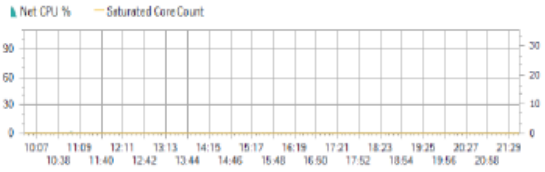

No GPU

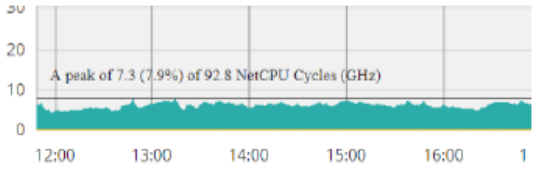

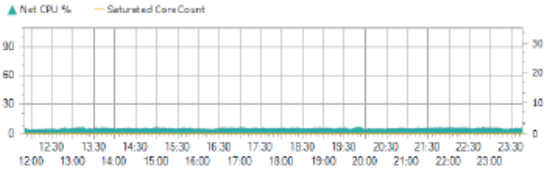

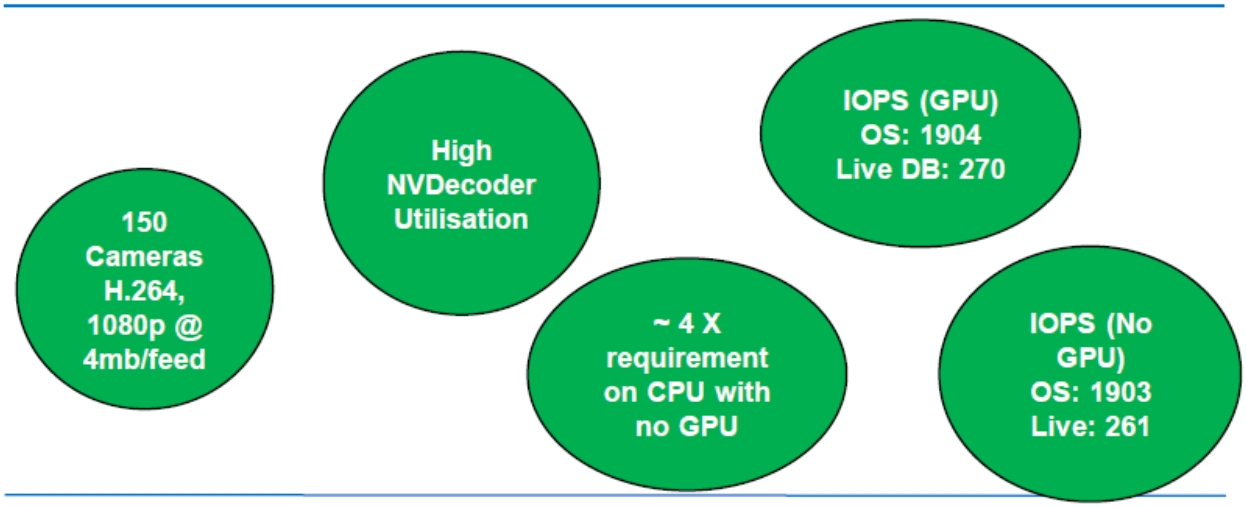

5.5 Conclusion

By introducing T4 GPU accelerator capability into the environment significantly enhances and increase the performance of the Chassis based solution whilst ensuring that the efficiency and best TCO outcome can be maintained. From the results we can see that with the NV-Decoder on the T4 working there is at least a greater than 4 x the offload from the CPU which is a considerable number.

We can also deduce that the recording servers could easily scale to 2 or 3 times of cameras access them at a single point in time.

The CUDA Cores were also being leveraged during the testing as Motion detect was enabled on all feeds. Memory utilization was always < 32GB and IOPS consistent between both series of testing.

6 Sample Design Architecture

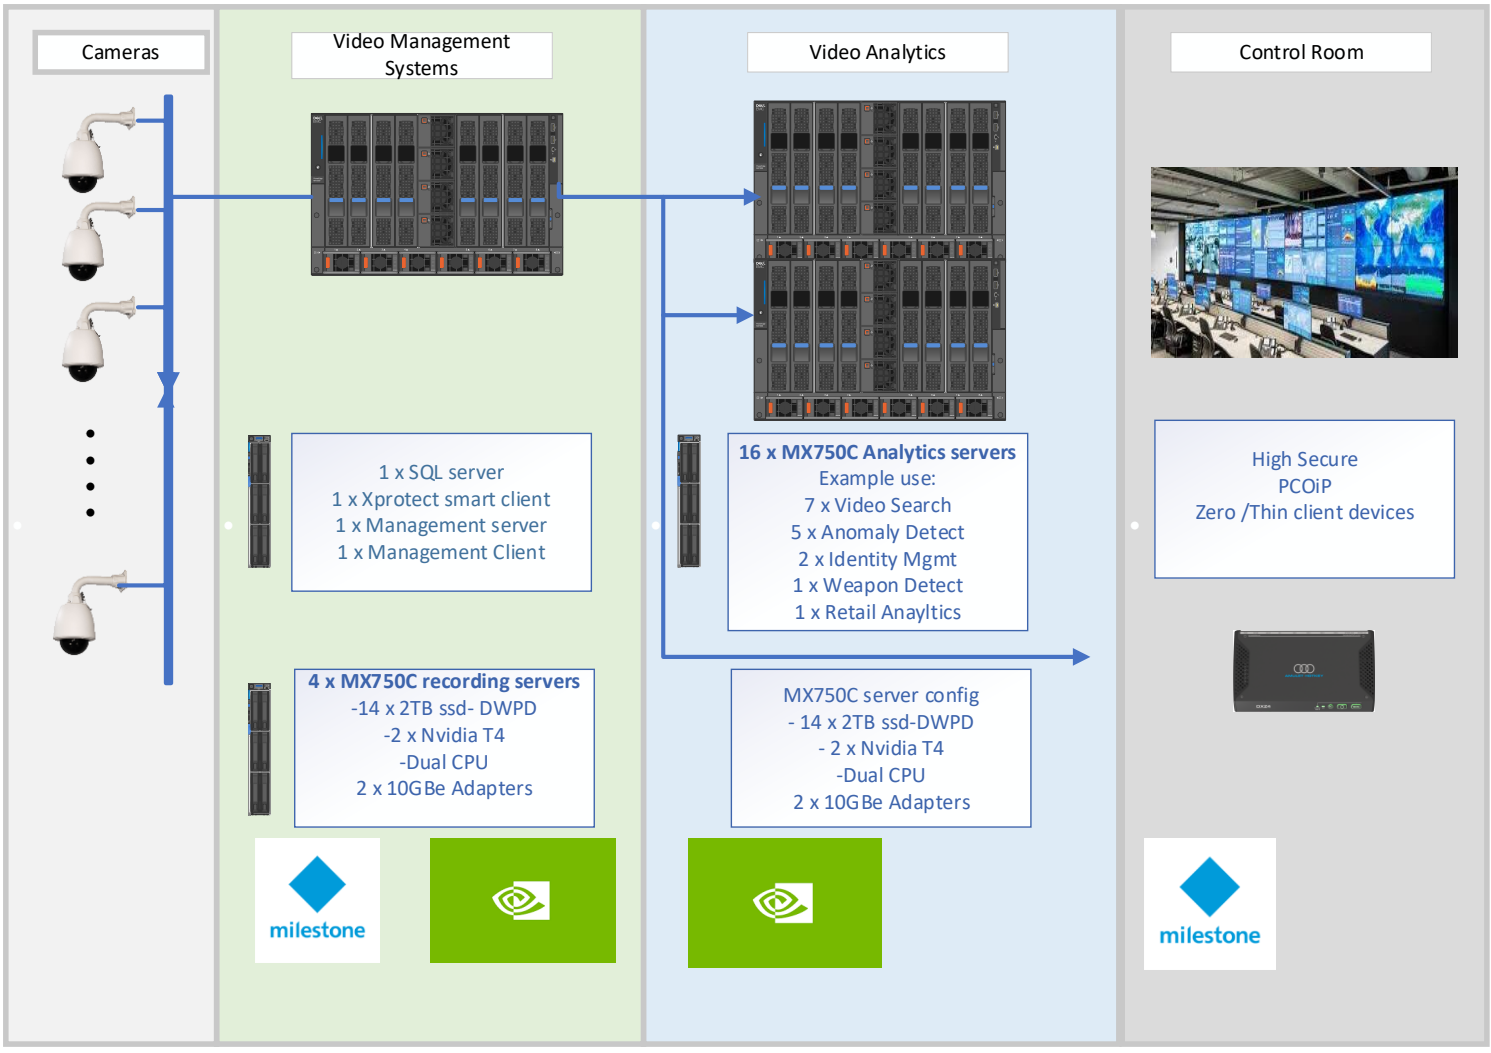

The Above solution show a typical use case example where not only is Milestone Xprotect being utilised but also, we have incorporated leading ISV analytic software stack to deliver a complete comprehensive end to end solution. In the example above you could easily have 1000+ cameras deployed, with external storage attached the data can be analysed in real time with the outputs being securely sent to a reginal command and control location and the data being displayed on large smart walls. Any incident can be captured highlighted and acted upon in real time.

Summary

Opportunity

Few companies have looked at the use of Chassis Based VMS offerings due to the Min size and scale at which is needed to make the solution cost effective. However if you start considering all the other components wrapped around the basic VMS solution, such as dedicated DB server, management servers, client servers, analytics (adding in additional capability offered by unique ISV companies) this then make the chassis based solution much more feasible. Very cost effective and better TCO returns.

Unique expertise

Amulet Hotkey's experience in optimized GPU solutions combined with innovative product design and Tier-1 relationships with vendors such as Dell Technologies and NVIDIAprovides us a unique position to design the optimal Computer vision solution.

Adaptation with a competitive edge

At a time when the world is undergoing major change driven by unstable financial markets, Data security challenges and lockdowns resulting from a global pandemic, the need for secure and flexible computer vision solutions which can match the changes each company faces as it tries to enter into a post Pandemic world ensuring the safety of all its employees, customers means having an infrastructure in place that is scalable flexible to meet these demands is greater than ever. Successfully being able to offer this flexibility whilst providing 'business as usual' to everyone and the reassurance that all is being done to ensure a safe working environment will give your organization a significant competitive edge and enable your key users to continue working in any situation without compromise.

Adapting to, and managing change is key to success.

Contact us today to learn more.

7 Resources and links

- General support information, documentation and driver/firmware files:

https://resources.amulethotkey.com/resources

- CoreModule GPU acceleration pages on Amulet Hotkey support site:

https://www.amulethotkey.com/main_products/coremodule-mx/

- Amulet Hotkey CoreModule datasheet

https://resources.amulethotkey.com/download/CoreModule_Datasheet_DS-MX7K-0001.pdf

- Dell EMC's landing page for MX7000 Modular chassis, with technical specifications, use cases, support information and deep dive technical articles

  https://www.dell.com/en-uk/work/shop/productdetailstxn/poweredge-mx7000

- NVIDIA site for general GPU virtualization articles

https://www.nvidia.com/en-gb/data-center/virtualization/