

The Dell Precision T5820 tower as high performance Smart Client leveraging NVIDIA GPU hardware acceleration

Dell IP Video Platform Design and Calibration Lab

June 2018

H17414

Reference Architecture

Dell EMC Solutions



Copyright © 2018 Dell Inc. or its subsidiaries. All rights reserved.

Published June 2018

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS-IS.” DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

- [Document overview](#)..... 4
- [Hardware overview](#)..... 5
- [Test setup](#)..... 6
- [Test results](#)..... 7
- [Findings and design considerations](#)..... 12

Document overview

This document provides an overview for using the Dell Precision T5820 tower as a high performance Smart Client that leverages NVIDIA GPU hardware acceleration.

The T5820 tower features high performance in a new, innovative, versatile, compact design, that also includes support for NVIDIA Quadro GPU adapters. It is an ideal platform for the surveillance industry. Milestone has also introduced support for NVIDIA GPU rendering with the 2018R1 release of Smart Client

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 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 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.

XProtect Smart Client

XProtect Smart Client is the main client for the VMS, offering a full set of advanced features and designed for daily use by dedicated operators.

XProtect Smart Client is designed to run remotely from the operator's computer. It supports multiscreen usage in full screen mode or in floating windows mode.

Hardware overview

Dell Precision T5820 tower

The Dell Precision T5820 tower is an ideal performance platform to leverage Milestone Smart Client.

- Intel single-socket XEON-based platform
- Enterprise-Grade platform, reliable for sustained, extreme workloads
- Multiple storage options
- Expandable up to 5 local 4 TB SAS 7200 RPM 12Gb/s 3.5" high density disks
- Ideal for exporting of data

Table 1 Dell Precision T5820 platform

System configuration	T5820
CPU model	Intel XEON W-2125 @ 4.01 GHz
RAM	2 x 8 GB DIMM
Disks	1 x NVMe PM981 500G
Graphic card	NVIDIA P2000, NVIDIA P4000, and NVIDIA P600
OS	Win10 pro 1709
Milestone XProtect release	2018 R1 2.1a build 7751
BIOS/Drivers	OS Update, + system updates with BIOS

NVIDIA Quadro

Add GPU acceleration by including a NVIDIA based GPU graphics adapter. On tests, we used the NVIDIA P2000, P4000, and P600 GPU adapters.

- Designed for enterprise-grade reliability and 24x7 operation
- Fully tested and qualified on Dell's workstation platforms
- Quad display ports for multi monitor support
- Long product life cycle
- H.264 and H.265 decoding, supported by Milestone
- Scalable to enable multi-GPU, multimonitor, and additional GPU decoders

Table 2 NVIDIA QUADRO specification

Adapter	P2000	P4000	P600
CUDA Parallel-Processing Cores	1024	1792	384
GPU Memory	5 GB GDDR5	8 GB GDDR5	2 GB GDDR5

Table 2 NVIDIA QUADRO specification (continued)

Adapter	P2000	P4000	P600
FP32 Performance	3.0 TFLOPS	5.3 TFLOPS	—
Max Power Consumption	75 W	105 W	40 W
Graphics Bus	PCI Express 3.0 x 16	PCI Express 3.0 x 16	CI Express 3.0 x 16
Display Connectors	DP 1.4 (4)	DP 1.4 (4) Optional Stereo (1)	4x mDP 1.4
Form Factor	4.4" H x 7.9" L Single Slot	4.4" H x 9.5" L Single Slot	2.713" H x 5.7" L, Single Slot, Low Profile

NVIDIA cards v3.0 and newer ([codename Kepler](#)) support hardware accelerated decoding of H.264 and JPEG streams in the XProtect Smart Client, including the following:

- GeForce GTX 660, GTX 670, GTX 780, GTX 680M, GTX 750M, Quadro K5000, Quadro K4200, Quadro K4000, Quadro K2200, K3100M, Quadro K4000M, and Quadro K5000M.

NVIDIA cards v6.0 ([codename Pascal](#)) and newer support hardware accelerated decoding of H.265 streams in the XProtect Smart Client, including the following:

- GeForce TITAN Xp, Titan X, GTX 1080, GTX 1070, GTX 1060, GTX 1050Ti, GTX 1050, GT 1030, Quadro P6000, Quadro P5000, Quadro P4000, Quadro P2000, Quadro P1000, Quadro P600, Quadro P400, Quadro P5000 (Mobile), Quadro P4000 (Mobile), and Quadro P3000 (Mobile).

For more information NVIDIA hardware acceleration support, see [Using hardware acceleration \(NVIDIA, Intel\) in XProtect Smart Client 2018 R1 and newer versions](#).

Test setup

The test environment includes a dedicated Milestone VMS server, Milestone XProtect Smart Client 2018 R1 (build 7751) running on Microsoft Windows 2010 pro (1709), and the remote focus client system for test. The server is running an emulation environment that can be calibrated to induce precise scaling to the focus test platform. In this case, the video streams are calibrated to provide a 4.4 Mb/s throughput that comprises a H.264, 1080p, feed at 30 fps.

We leveraged the Milestone hardware diagnostics within the smart client to monitor the frame performance. We also compared live streaming versus playback.

When the optimum level is reached, the Dell LiveOptics tool gathers the system performance information over a four-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, the Milestone Performance and Microsoft Performance (Perfmon) Resource Monitor tools monitor performance and frame loss.

We used the `NVIDIA_SMI.exe`, Microsoft Task Manager, and Perfmon to monitor the system and GPU performance.

Test criteria

Sustained frame drops on the playback feeds were the criteria for stopping the testing. This was in line with the maximum load on the GPU, and CPU.

Data points gathered

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

Microsoft system resource monitor

- CPU % utilization
- GPU % utilization
- Memory usage
- GPU memory utilization

LiveOptics 4 hour analysis at max performance

- Application diagnostics measuring
- Frames per second
- Codec
- Resolution
- Hardware acceleration
- Frames per second, for playback and live view

Test results

Quadro P2000 - Single Display running 1080p 30FPS Feeds

Figure 1 P2000 GPU performance

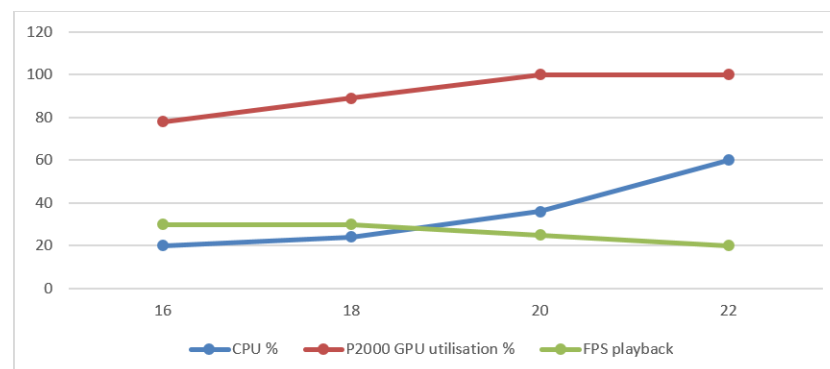


Table 3 P2000 GPU results

No. of feeds	CPU %	GPU Utilization %	FPS playback
		P2000	
16	20	78	30
18	24	89	30
20	36	100	25
22	60	100	20

Quadro P4000 - Single Display running 1080p 30FPS Feeds

The P4000 adapter has a similar peak feed performance to the P2000; however it makes less of a demand on the CPU.

Figure 2 P4000 GPU performance

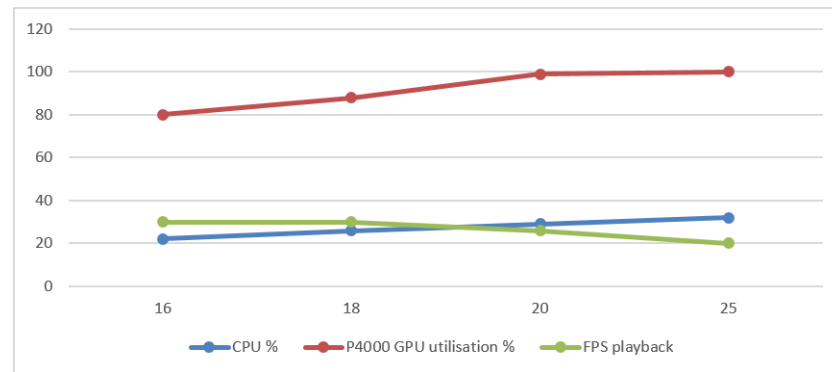


Table 4 P4000 GPU results

No. of feeds	CPU %	GPU Utilization %	FPS playback
		P4000	
16	22	80	30
18	26	88	30
20	29	99	26
25	32	100	20

Quadro P600 - Single Display running 1080p 30FPS Feeds

Figure 3 P600 GPU performance

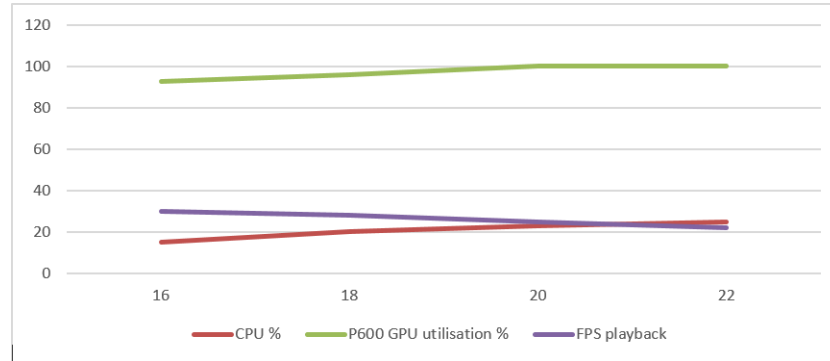


Table 5 P600 GPU results

No. of feeds	CPU %	GPU Utilization %	FPS playback
		P600	
16	15	93	30
18	20	96	28
20	23	100	25
22	25	100	22

Dual GPU Adapter

Adding additional GPU adapters leverages both in parallel performance, effectively doubling the quantity of feeds being displayed. It also provides additional output ports if required.

Dual GPU performance

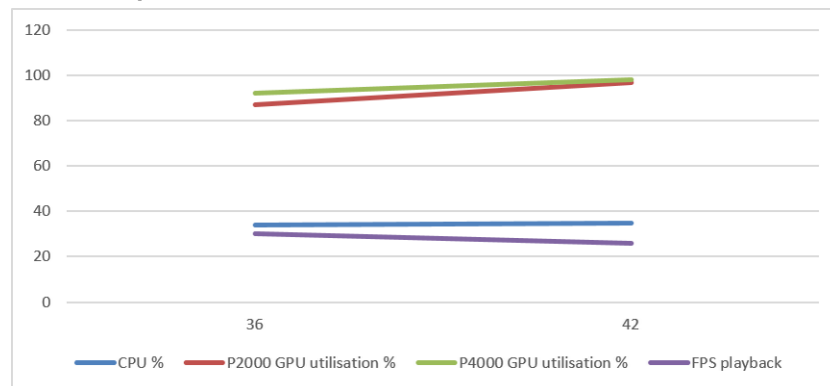


Table 6 Dual GPU results

No. of feeds	CPU %	GPU Utilization %		FPS playback
		P2000	P4000	
36	34	87	92	30
42	35	97	98	26

Multiple Monitors, Multiple instances of the client per screen

Additional instances of the Smart Client place a higher demand on the CPU. The overall quantity of feeds being displayed saturates at approximately 20 percent less than if utilizing a single display.

The performance of both GPU's are being leveraged in parallel.

Figure 4 Multi-monitor / Dual GPU performance

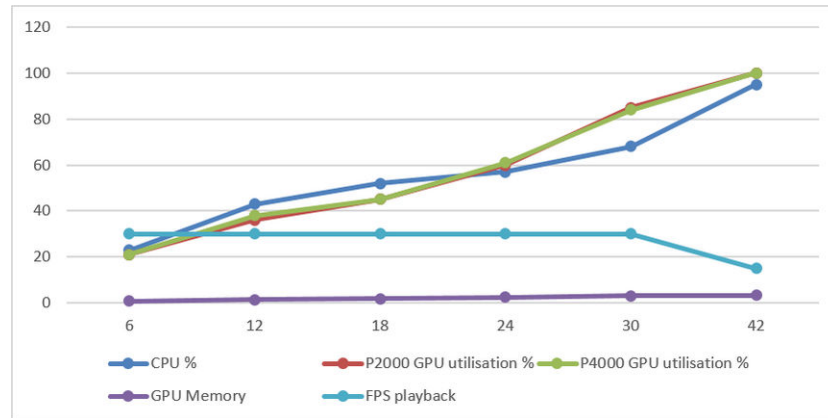
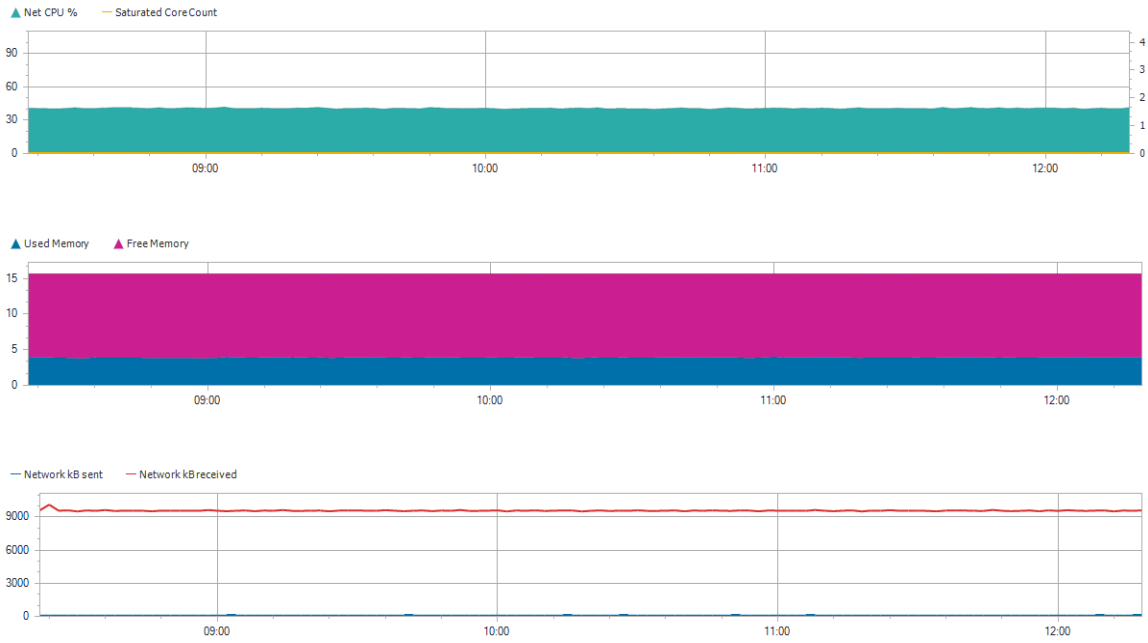


Table 7 Multi-monitor/Dual GPU results

Instances of Client	No. of displays	No. per screen	No. of feeds	CPU %	GPU Utilization %		GPU Memory	FPS playback
					P2000	P4000		
1	5	6	6	23	21	21	0.7	30
2	5	6	12	43	36	38	1.4	30
3	5	6	18	52	45	45	1.8	30
4	5	6	24	57	60	61	2.5	30
5	5	6	30	68	85	84	3.1	30
5	5	6+(1*12)	42	95	100	100	3.2	15

System performance over a 4 hour period at multi-monitor (Data from LiveOptics)



Findings and design considerations

- Performance is greatly improved by leveraging GPU hardware acceleration.
- The Smart Client leverages multiple GPU adapters when present, operating them in parallel.
- The feeds resolution and frame rate have a linear demand on the playback decoding.
- The best maximum performance of the GPU adapter is at approximately 80 percent before it increases its reliance on the CPU, and starts to decrease the Frame Per Second playback.
- Exporting of video feeds has a minimal impact on the GPU and CPU (approximately 3 percent). Tested with 10 feeds exporting over one hour.
- The Milestone Smart Client primarily leverages the NVDEC decoder chip on the adapters for accelerating the decoding of the video streams. It is important to choose the Pascal chipset or above to include support for H.264, H.265 video streams.
- The Quadro P600 adapter is an excellent choice for a single display that can handle up to 16 HD streams.
- If the streams are HD and above, such as 4K, the P2000 adapter performs better.
- Also consideration on the playback screen is important, especially when playback of high-resolution images is required.
- On tests, we leveraged multiple display types. The Dell P4317Q monitor, with a native resolution of 3840x2160, is an ideal choice for a physically large display at 43" with high-resolution.