Total Cost of Ownership (TCO) Validation Study: *Dell EMC HPC Ready Architecture for AI and Data Analytics vs. Dell EMC HPC, AI and Data Analytics Reference Architectures*

Silverton Consulting, Inc. StorInt™ Briefing

Introduction

Dell Technologies has long provided state-of-the-art servers, storage and networking solutions that feature latest generation hardware componentry from a number of vendors. This equipment has been deployed in the enterprise and other market spaces across a broad range of verticals and workloads, including enterprise online transaction processing (OLTP); high-performance computing (HPC) and scientific computing; artificial intelligence (AI), machine learning (ML) and deep learning (DL); and data analytics, big data and data warehouses.

The computing services used for HPC, AI/ML/DL and Data Analytics have historically been deployed across separate, distinct hardware and software infrastructure. HPC scientific computing workloads were optimized for high-performance, dual-precision, floating-point arithmetic CPUs; high-memory bandwidth, high-speed interconnects; and high-capacity/high-sequential bandwidth data storage. AI/ML/DL processing relied on multiple graphic processing unit (GPU) or field-programmable gate array (FPGA) processing cores and high-capacity/high-sequential bandwidth data storage. Data Analytics depended on high numbers of commodity server cores with high-capacity, direct-access server storage.

Today, businesses are starting to turn to HPC-like modeling and simulation to optimize products/services, operations and sales/marketing activities. Historically, this role had been filled by AI and Data Analytics processing. Moreover, the data used by these three workloads can have a lot of similarities. In some cases, the same data load (augmented, filtered and cleansed) has been used by all three. Using separate infrastructure for these workloads often required IT time and effort to move or replicate data from one system to another for processing the workload.

Recently, HPC software is being revised to take advantage of mixed-precision floating-point hardware. AI/ML/DL processing and Data Analytics computations are also moving to smaller precision floating point. With the emergence of mixed-precision floating-point hardware and multi-precision HPC, AI/DL and Data Analytics software applications, all of these processing activities can now use the same processing hardware and data storage.

We see evidence of this HPC-AI/DL-Data Analytics convergence in recent supercomputers being deployed, such as the ENI HPC5 with total peak power of 70 PFLOPS (quadrillion floating-point operations per second) of computing across ~ 1,820 Dell EMC PowerEdge C4140 servers, each with two Intel Gold 6252 24-core processors and four NVIDIA V100 GPU accelerators for AI and HPC computing activities.[1]

[1] Please see https://www.eni.com/en-IT/media/press-release/2019/10/eni-announces-plans-for-a-new-supercomputer-at-the-green-data-center.html

Along those lines, Dell Technologies has recently introduced a new **HPC Ready Architecture for AI and Data Analytics**, which combines HPC-AI/DL-Data Analytics in one infrastructure, that customers can order today and deploy tomorrow.

This paper was funded by Dell Technologies to validate the ***TCO [Total Cost of Ownership] Analysis—HPC Ready Architecture for AI and Data Analytics***.[2] In order to confirm that solution brief's TCO analysis, we requested and received updated formal price quotes for all of the systems discussed in the original brief. Our costs differ somewhat from those in the original brief; however, system prices often differ over time, primarily due to volatile commodity and non-commodity component pricing. [3]

Below, we discuss our TCO comparison of Dell Technologies HPC Ready Architecture for AI and Data Analytics (hereafter called the **HPC Unified Architecture**) vs. three separate Dell Technologies reference architectures for HPC, AI and Data Analytics. This comparison affirms a significant **(for our analysis, ~$1.5M less)** three-year cost advantage of using the HPC Unified Architecture over the three separate systems.

Dell EMC HPC Ready Architecture for AI and Data Analytics

The HPC Unified Architecture integrates all of the compute performance, accelerators and other hardware in a single infrastructure needed to run HPC-AI-Data Analytics workloads together.

The HPC Unified Architecture uses the following three Dell EMC servers:

- **PowerEdge R740xd servers**, which are 2U accelerated servers that offer up to 24 drives/NVMe SSDs of onboard direct storage; several dual-socket, high-performing Intel® processor options; high-capacity DRAM memory configurations; and up to two PCIe NVIDIA® Tesla V100 GPUs, NVIDIA T4 GPUs or Intel Programmable Acceleration Cards with Intel Arria® or Stratix® FPGAs.
- **PowerEdge C4140 (configuration M) servers**, which are dense, 1U accelerated servers that support diskless/SSD-less systems with several dual-socket, high-performing Intel® processor options; high-capacity DRAM memory configurations; and up to four NVIDIA Tesla V100 GPUs with PCIe or optional NVIDIA **NVLink®** technology fabrics.

[2] Please see https://infohub.delltechnologies.com/section-assets/h18136-tco-analysis-dell-emc-hpc-ra-for-ai-da-sb as of 25 Feb 2020.
[3] Price quotes available upon request.

- **DSS 8440 servers**, which are dense, 4U accelerated servers with several dual-socket, high-performing Intel® processor options. A single server supports up to 10 NVIDIA Tesla V100 GPUs, 16 NVIDIA T4 GPUs or special-purpose AI/DL intelligence processing units (IPUs).

The HPC Unified Architecture also supports two different interconnect networking fabrics, one based on 100 Gbps InfiniBand (using Mellanox® hardware) and the other based on 25 Gbps Ethernet (using Dell EMC hardware).

Supporting all three workloads requires HPC and Data Analytics compute nodes, AI/DL training and inferencing GPU nodes and a cluster management node. In the HPC Unified Architecture configuration selected for the TCO analysis, we used the InfiniBand network option and the following server selections:

- **1 PowerEdge R740xd server** with dual-socket Intel Xeon Gold 6248 processors (20C/40T, 2.5Ghz), 384GB of DRAM and 24 3.84TB SATA SSDs for the management or cluster head node.
- **2 PowerEdge R740xd servers**, each with dual-socket Intel Xeon Gold 6248 processors, 384GB of DRAM, two 960GB SATA SSDs and one NVIDIA Tesla V100S GPU for compute and GPU nodes.
- **2 PowerEdge C4140 servers**, each with dual-socket Intel Xeon Gold 6248 processors, 384GB of DRAM, no disks/no SSDs and four NVIDIA Tesla V100 GPUs for compute and GPU nodes.
- **1 DSS 8440 server** with dual-socket Intel Xeon Gold 6248 processors, 384GB of DRAM, one 960GB SATA SSD and eight NVIDIA Tesla V100 GPUs for compute and GPU nodes.

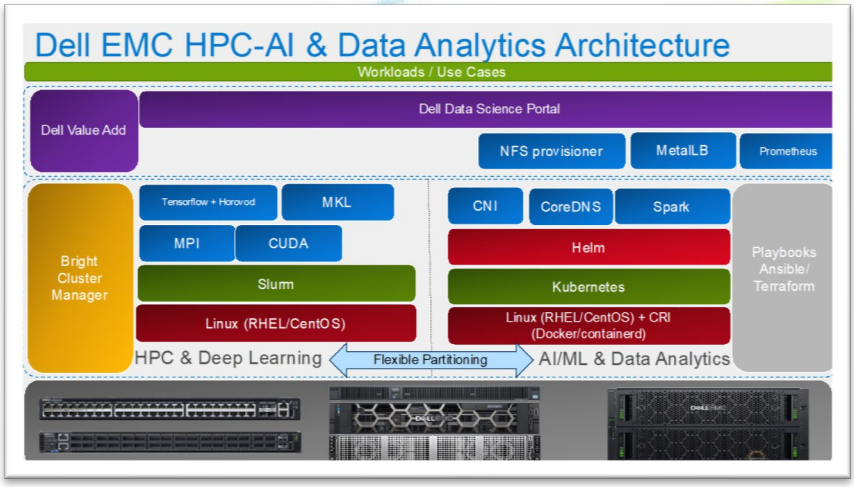

For the HPC Unified Architecture solution, we administer the cluster of servers with **Bright Cluster Manager®** software, which is used to allocate/deallocate (schedule/un-schedule) nodes to HPC, AI/DL or Data Analytics processing activities (e.g., **Slurm®** nodes for HPC batch work or **Kubernetes/Docker** nodes for AI/DL or Data Analytics activities). Dell Technologies supplies **Ansible** playbooks to reconfigure cluster nodes to be Slurm or Kubernetes/Docker nodes. Kubernetes/Docker container systems support several AI frameworks for AI/DL and several analysis solutions for Data Analytics.

Dell Technologies reference architectures for dedicated HPC, AI and Data Analytics

Below we describe the three separate reference architecture systems used to support dedicated HPC, AI and Data Analytics processing environments.

Dell EMC HPC Reference Architecture

The HPC Reference Architecture offers a wide selection of Dell EMC servers and other equipment depending on your industry vertical, including the PowerEdge R740xd and C4140 servers described above, as well as the following:

- **PowerEdge R640 servers**, 2U servers that support several dual-socket, high-performance Intel processor options, along with large memory configurations, with up to 10x 2.5" SAS/SATA disk or SSD drives or 8x NVMe SSDs plus two more 2.5" SAS/SATA disk or SSD drives.
- **PowerEdge C6420 servers**, density-optimized 2U chassis with four servers that support several dual-socket, high-performance Intel processor options and medium memory configurations with up to 24x 2.5" SAS/SATA hard drives or SSDs. The C6420 servers also offer **Intel® Omni-Path®** fabrics and **liquid cooling** options.

All Dell EMC HPC Reference Architectures support both Ethernet and InfiniBand networking fabrics options.

In the HPC Reference Architecture configuration used for our TCO analysis, we selected the InfiniBand network option and the following servers:

- **1 PowerEdge R640 server** with dual-socket Intel Xeon Bronze 3106 processor (8C/8T, 1.7Ghz), 192GB of DRAM and one 960GB SATA SSD for the management or cluster head node.
- **8 PowerEdge C6420 servers**, each with dual-socket Intel Xeon Gold 6242 processors (16C/32T, 2.8Ghz), 192GB of DRAM and two 480GB SATA SSDs for compute nodes.

Dell EMC AI Reference Architecture

The Dell EMC AI Reference Architecture supports **Deep Learning with Intel** and **Deep Learning with NVIDIA**:

- **Deep Learning with Intel** utilizes a R740xd server as the head or management node and up to 16 PowerEdge C6420 servers (described above) with Ethernet and a Kubernetes software stack with AI/ML/DL container workloads.
- **Deep Learning with NVIDIA** uses the R740xd as the head node and up to four PowerEdge C4140 servers (described earlier). Each server could include up to four NVIDIA Tesla V100 GPUs. Recall that the C4140 supports NVIDIA

NVLink as well as PCIe access for GPUs to support any AI/ML/DL workload requirements.

The AI Reference Architecture supports Ethernet or InfiniBand fabrics for networking node interconnect.

In the AI Reference Architecture configuration used for our TCO analysis, we selected InfiniBand networking and the following servers:

- **1 PowerEdge R740xd server** with dual-socket Intel Xeon Gold 6248 processors, 192GB of DRAM and 12 12TB SATA nearline disks for the management or cluster head node.
- **4 PowerEdge C4140 servers**, each with dual-socket Intel Xeon Gold 6248 processors, 192GB of DRAM, no disk/SSD and four NVIDIA Tesla V100 GPUs for AI training/inferencing nodes.

Dell EMC Data Analytics Reference Architecture

The Dell EMC Data Analytics Reference Architecture supplies dedicated environments to support three architecture options: **Hadoop**, **Big Data as a Service (BDaaS)** and **real-time data streaming**.

- The **Hadoop architecture** uses multiple (up to 288) R640 or R740xd servers with up to 64TB of direct-attached storage (DAS) per node, with Ethernet or InfiniBand networking running either Cloudera Enterprise Hadoop or Hortonworks Data Platform (HDP) Hadoop software.
- The **BDaaS architecture** uses PowerEdge R640 servers as an administrator or gateway node and PowerEdge R740xd servers as either Intel-based or GPU-based accelerator nodes, with Ethernet or InfiniBand networking running BlueData EPIC or custom images for Cloudera Hadoop, Hortonworks Hadoop, Cassandra NoSQL, Spark or in-memory GPU services.
- The **real-time data streaming architecture** uses multiple PowerEdge R640 servers with up to 12TB of DAS disk or 9.6TB of NVMe SSD storage per node, using Ethernet fabric running Confluent Control Center, Confluent Platforms, Confluent KSQL and Kafka Brokers.

In the Data Analytics Reference Architecture configuration used for our TCO analysis, we selected InfiniBand networking and **five PowerEdge R640 servers**, each with dual-socket Intel Xeon Gold 6248 processors; 384GB of DRAM; and two 960GB SATA SSDs, two 1.92TB SATA SSDs and two 480GB SATA SSDs for DAS.

Three-year TCO comparison: HPC Unified vs. three separate architectures

We modeled the acquisition and running costs for the four configurations described above over the course of three years. The specific costs modeled for the systems include **hardware and software acquisition costs**; **staffing costs**; and **management costs**, **power and cooling costs**.

Dell Technologies supplied all hardware, storage and networking pricing. Software costs include Dell software plus a mixture of open source and non-open source solutions. System acquisition costs were modeled for the four configurations using Dell official price quotes (as noted earlier, all price quotes are available on request). Hardware and software costs are included only in year 1 and are not included in years 2 and 3 as no new hardware or software was added during those years.

Staffing costs are based on a **$67K/year admin salary** used for all three years. We assumed that the staffing remains flat for the three-year period for all systems and used a three-admin headcount for each of the three separate systems and five admins for the HPC Unified system. The increased headcount for HPC Unified is required to manage the additional server hardware.

Management costs and power and cooling costs are a function of system size and are equal to **30% of the acquisition costs**. Again, no hardware or software was added in years 2 and 3, so these costs remain flat for those years.

In the analysis that follows, we compare the total costs of deploying all three of the separate, distinct HPC, AI and Data Analytics systems against the cost of a single HPC Unified system. All of our TCO charts show the total costs for the three separate configurations in comparison with the cost of the HPC Unified Architecture.

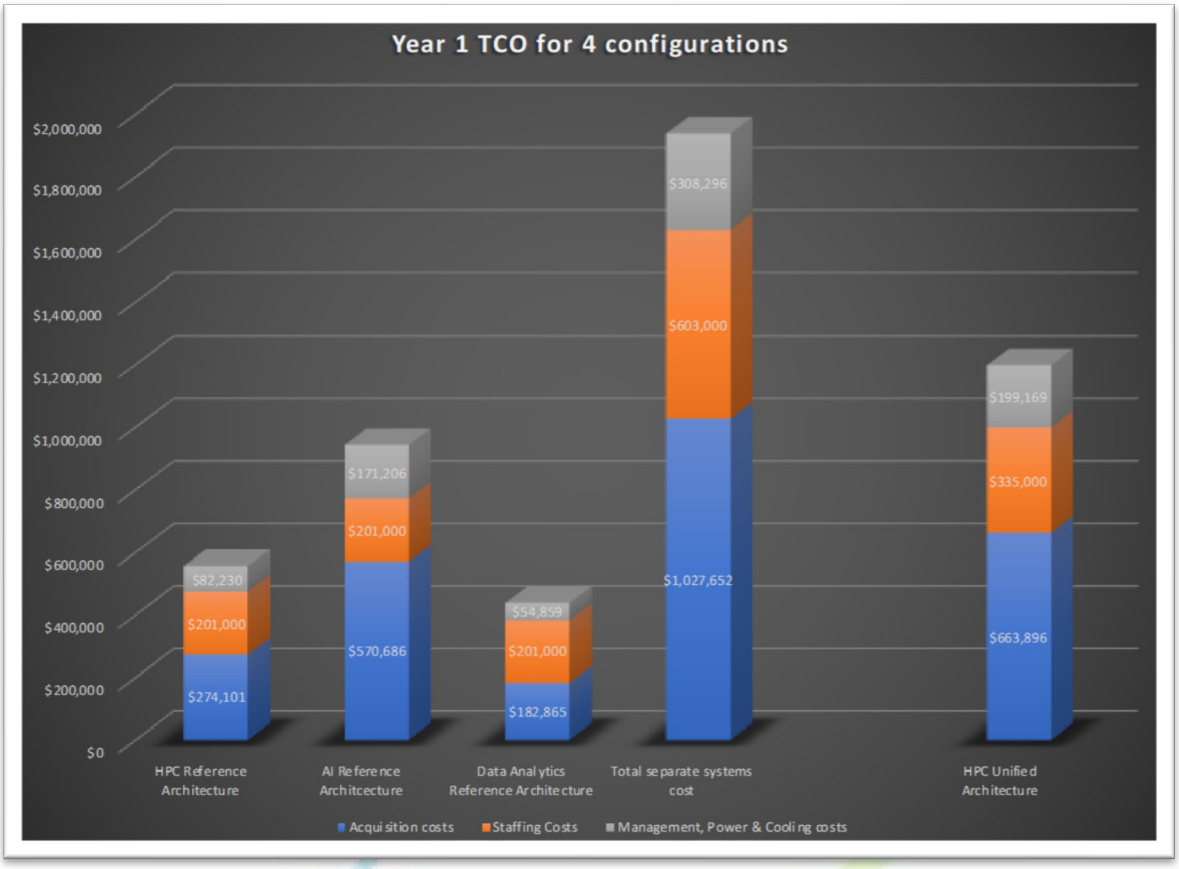

Total costs for Year 1 for the four configurations are as follows:

- **HPC Reference Architecture** hardware and software acquisition costs for the single R640 head node and the eight C6420 servers with InfiniBand fabric is **$274.1K**. We used a three-admin headcount to come up with staffing costs of **$201.0K**. Management, power and cooling costs were modeled at 30% of acquisition costs, totaling **$82.2K**. The total year 1 cost for the HPC Reference Architecture configuration is **$557.3K**.
- **AI Reference Architecture** hardware and software acquisition costs for the single R740xd head node and the four C4140 compute nodes with four GPUs each with InfiniBand fabric is **$570.7K**. We used a three-admin headcount to come up with staffing costs of **$201.0K**. Management, power and cooling costs were modeled at 30% of acquisition costs, totaling **$171.2K**. The total year 1 cost for the AI Reference Architecture configuration is **$942.9K.**
- **Data Analytics Reference Architecture** hardware and software acquisition costs for the five R640 compute nodes with InfiniBand fabric is **$182.9K**. We used a three-admin headcount to come up with staffing costs of **$201.0K**. Management, power and cooling costs were modeled at 30% of acquisition

costs, totaling **$54.9K**. The total year 1 cost for the Data Analytics Reference Architecture configuration is **$438.7K**.

- The **three separate configurations** include acquisition costs totaling **$1027.7K**; staffing costs totaling **$603.0K**; and management, power and cooling costs totaling **$308.3K**. The total year 1 cost is **$1,938.9K** for the three systems.
- **HPC Unified Architecture** hardware and software acquisition costs for the single R740xd head node, the two R740xd with one GPU compute node, the two C4140s with four GPUs each per node and the single DSS 8440 with eight GPUs, with InfiniBand fabric total **$663.9K**. Staffing costs (five-admin headcount) total **$335.0K**. Management, power and cooling costs were modeled at 30% of acquisition costs, totaling **$199.2K**. The total year 1 cost for the HPC Unified configuration is **$1,198.1K**.

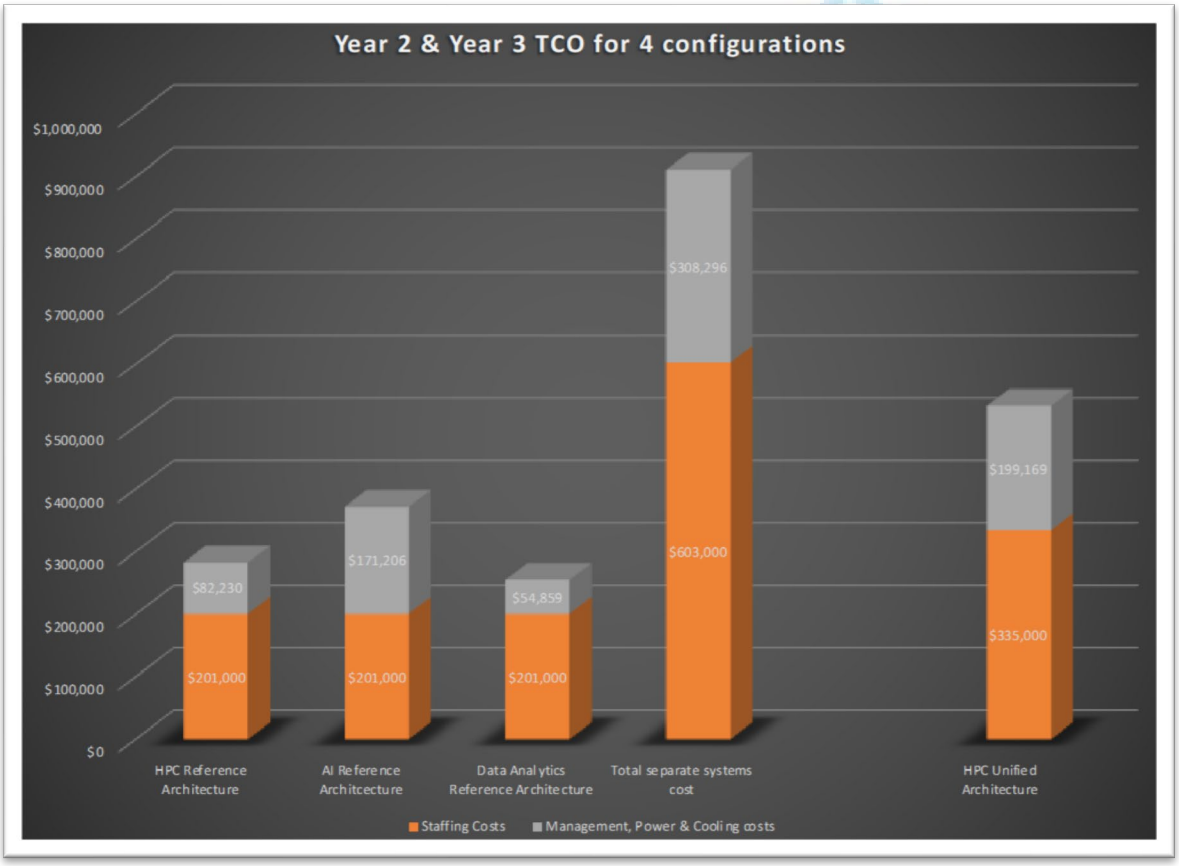

Total costs for Years 2 and 3 for the four configurations are the same over the two years and are as follows:

- **HPC Reference Architecture** staffing (three-admin headcount) costs **$201.0K**, and management, power and cooling cost **$82.2K**. The total year 2 and year 3 cost for the HPC Reference Architecture configuration is **$283.2K**.

- **AI Reference Architecture** staffing (three-admin headcount) costs **$201.0K**, and management, power and cooling cost **$171.2K**. The total year 2 and year 3 cost for the AI Reference Architecture configuration is **$372.2K**.
- **Data Analytics Reference Architecture** staffing (three-admin headcount) costs **$201.0K**, and management, power and cooling cost **$54.9K**. The total year 2 and year 3 cost for the Data Analytics Reference Architecture configuration is **$255.9K**.
- The **three separate configurations** include a nine-admin headcount totaling **$603.0K**. Management, power and cooling cost **$308.3K**. The total year 2 and year 3 cost for the three systems is **$911.3K**.
- The **HPC Unified Architecture** five-admin headcount costs **$335.0K**, and management, power and cooling cost **$199.2K**. The total year 2 and year 3 cost for the HPC Unified configuration is **$534.2K**.

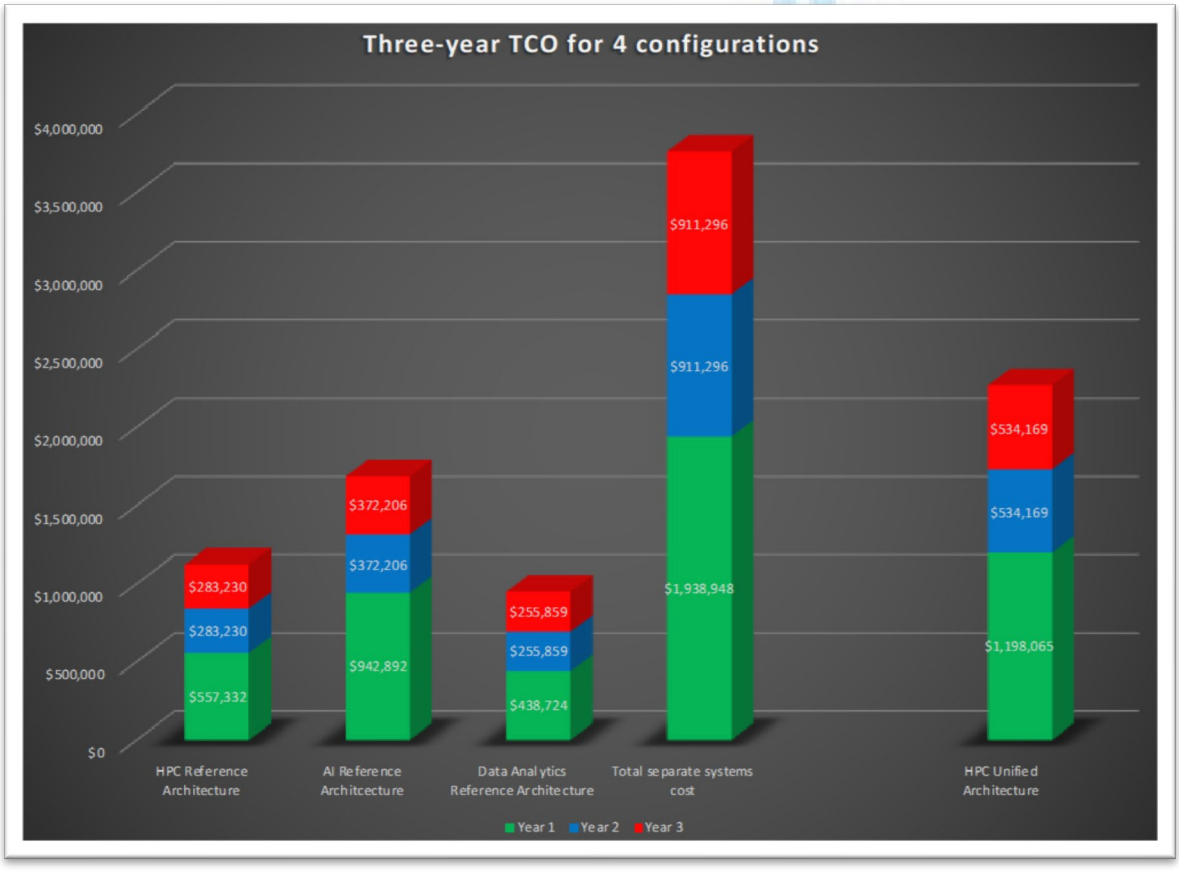

The three-year total costs for the systems are as follows:

- **HPC Reference Architecture** year 1, 2 and 3 costs are **$557.3K**, **$283.2K**, and **$283.2K**, respectively, for a total three-year cost of **$1,123.8K**.
- **AI Reference Architecture** year 1, 2 and 3 costs are **$942.9K**, **$372.2K** and **$372.2K**, respectively, for a total three-year cost of **$1,687.3K**.

- **Data Analytics Reference Architecture** year 1, 2 and 3 costs are **$438.7K**, **$255.9K**, and **$255.9K**, respectively, for a total three-year cost of **$950.4K**.

The total costs for the three separate systems across the three-year period are **$1,938.9K** for year 1, **$911.3K** for year 2 and **$911.3K** for year 3, for a total three-year cost of **$3,761.5K.**

The HPC Unified Reference Architecture total costs across the three-year period are **$1,198.1K** for year 1, **$534.2K** for year 2 and **$534.2K** for year 3, for a total three-year cost of **$2,226.4K**, **~40% (or $1.495.1K) less** than the costs of the three separate systems.

Of course, actual costs may vary. We also have not modeled additional costs such as inflation, fully burdened salaries, building floor space, etc., which may impact the TCO of one or more of the above configurations. In addition, as discussed earlier, hardware component prices change over time, so future system pricing may differ from the pricing used here. However, given the magnitude of cost savings shown above, we firmly believe any cost comparison would show a significant cost savings resulting from the deployment of the HPC Unified Architecture over the deployment of three separate systems.

Performance equivalence of HPC Unified vs. three separate architectures

Although our validation study was not intended to show performance equivalence, we offer the considerations shown below.

	CPU Cores	Memory (GB)	GPUs	Disk Capacity (GB)	SSD Capacity (GB)
HPC Architecture	272	1536	-	-	7,680
AI Architecture	200	960	16	144,000	-
Data Analytics Architecture	200	1920	-	-	33,600
Totals for 3 Systems	**672**	**4416**	**16**	**144,000**	**41,280**
HPC Unified Architecture	**240**	**2304**	**18**	**-**	**96,960**

The table includes totals for the CPU cores, memory size and GPUs, as well as disk and SSD capacity for all of the system servers. For CPU cores, memory and disk capacity, the HPC Unified system uses less hardware than the three separate systems. The following also holds true:

1. While the three separate systems have three separate stacks, the HPC Unified has a single (~1.3 for Slurm and Kubernetes) software stack, which reduces the CPU cores, memory and storage capacity required to support the workloads together.

2. The HPC Unified uses more GPU cores to perform HPC processing, which substantially reduces the CPU cores and memory required to support HPC workloads.
3. The HPC Unified needs only one copy of data, which reduces the storage capacity required to support the workloads.

On the other hand, the HPC Unified system needs to have sufficient capacity to match the maximum capacity of any of the three workloads. The AI system has 144TB of (nearline) disk capacity, while the HPC Unified has only 97TB of SSD.

In the HPC Unified system's defense, 97TB of SSDs are much more expensive than 144TB of nearline disk. To match the 144TB capacity, we could configure the HPC Unified system with 96TB of nearline disk and the remainder in SSD storage. However, this would reduce the cost of the HPC Unified system even more and it wouldn't perform as well as having all 97TB of SSD.

In other words, although the HPC Unified CPU core counts, memory size and storage capacity are not equal to those of the three separate systems we believe the numbers are close enough not to matter.

Summary

The total cost, including hardware and software acquisition, staffing, management, power and cooling costs, for the Dell Technologies HPC Ready Architecture for AI and Data Analytics costs **significantly less** than deploying three separate Dell Technologies configurations for HPC, AI and Data Analytics.

For our TCO analysis, we modeled system acquisition costs (Dell price quotes), staffing and management, power and cooling (at 30% of acquisition) costs. The staffing costs all use the same salary and the same headcount for the three separate systems and include a larger headcount for the HPC Unified Architecture to handle the increased hardware.

Although the three separate systems and the HPC Unified Architecture did not have identical CPU core counts, memory size and storage capacity, it is our firm belief that the HPC Unified Architecture can run the three workloads as well as the three separate systems.

Given this analysis, if your organization needs to support HPC, AI and Data Analytics workloads, it will make good business and financial sense to deploy the Dell EMC HPC Ready Architecture for AI and Data Analytics.