Hewlett Packard Enterprise

Technical White Paper

HPE CONVERGEDSYSTEM 300 FOR MICROSOFT ANALYTICS PLATFORM

Solution Architecture and Reference Guide for Microsoft Analytics Platform System, Cumulative Update 7.7 Release

CONTENTS

EXECUTIVE SUMMARY

Audience level: Executive, Technical, Sales, and PreSales.

HPE ConvergedSystem 300 for Microsoft Analytics Platform addresses the key needs of business insights through greater understanding of business data. It is pre-configured, pre-loaded with Microsoft® software and a purpose-built solution customized for Data Warehouse. As a proven Massively Parallel Processing (MPP) data warehouse solution, it can handle high concurrency and highly complex queries at scale.

The HPE ConvergedSystem 300 for Microsoft Analytics Platform can be easily and rapidly deployed to store and analyze all your data with high performance, concurrency and simplicity. All required components, including servers, storage, and networking, are integrated into a single converged system that is expandable to meet virtually any big data challenge.

HPE ConvergedSystem 300 for Microsoft Analytics Platform combines the best in breed Hewlett Packard Enterprise hardware solution with the software from Microsoft, along with world class support provided by both companies. The HPE ConvergedSystem 300 for Microsoft Analytics Platform allows you to start small and grow your converged system later as business needs change.

FIGURE 1. The solution scales from a "Quarter Rack" up to 7 racks

HPE ConvergedSystem 300 for Microsoft Analytics Platform is available as a system from as small as a "Quarter Rack" solution all the way up to 56 Active Scale Unit nodes or 64 total nodes, a total of up to 7 racks of hardware for a single HPE ConvergedSystem 300 for Microsoft Analytics Platform.

The solution offers up to 12.6 Petabytes (PB) of data storage using 6TB drives. Storage drive choices of 2TB or 6TB are available to help meet varying business requirements for data capacity.

Delivers scalability and ground-breaking performance

- HPE ConvergedSystem 300 for Microsoft Analytics Platform offers flexibility to scale to large amounts of data, up to 12.6 PB total.
- Designed with HPE servers and storage to deliver superb performance with all components balanced to reduce bottlenecks.
- The HPE ConvergedSystem 300 for Microsoft Analytics Platform has been designed with a dense server and storage configuration, and designed from the ground up.

Achieves time-to-application value with lowered risk

Hewlett Packard Enterprise and Microsoft have jointly developed, configured, tested, tuned, and validated the HPE ConvergedSystem 300 for Microsoft Analytics Platform, resulting in reduced implementation risk and cost.

- Each HPE ConvergedSystem 300 for Microsoft Analytics Platform is delivered along with a unique collaborative support agreement between HPE and Microsoft. With a unified solution support experience, disruptions are minimized, and unnecessary downtime avoided.
- Included HPE Factory Express Deployment Service for the HPE ConvergedSystem 300 for Microsoft Analytics Platform facilitates a proper customer site assessment, factory integration, on-site installation services and end-to-end product management from ordering through go-live stage, which is performed to allow smooth installation and allow you to focus on business issues rather than infrastructure.
- Automated installation and validation tools, utilized at your site by a team of HPE and Microsoft engineers, provide consistent and error free deployment. All software is preinstalled, and the hardware components are preconfigured in the HPE factory.
- One of the innovations that HPE has developed is a unique set of tools designed to reduce operational costs. The HPE Support Pack includes tools to significantly simplify updates and several other maintenance tasks across this multi-node appliance. Collectively, these tools help ensure the system is running at optimal performance and are invoked with nothing more than a single mouse click, returning the benefits expected from your investment. This support pack ships with the HPE ConvergedSystem 300 for Microsoft Analytics Platform.

Target audience: The target audience for this document is chief information officers (CIOs), chief technology officers (CTOs), IT directors, IT architects, database administrators, business intelligence (BI) administrators, and others wishing to learn more about this solution.

This document assumes the reader has a basic understanding of several key data center technologies including, but not limited to servers, storage, networking, power, solution management, virtualization, and hypervisors.

Disclaimer: Products sold prior to the separation of Hewlett-Packard Company into Hewlett Packard Enterprise Company and HP Inc. on November 1, 2015 may have a product name and model number that differ from current models.

CURRENT RELEASE

The current release of the HPE ConvergedSystem 300 for Microsoft Analytics Platform is installed with Microsoft Analytics Platform System (APS), Cumulative Update (CU) 7.7 release. Microsoft Analytics Platform System hosts SQL Server Parallel Data Warehouse (PDW), which is the software that runs the massively parallel processing (MPP) data warehouse.

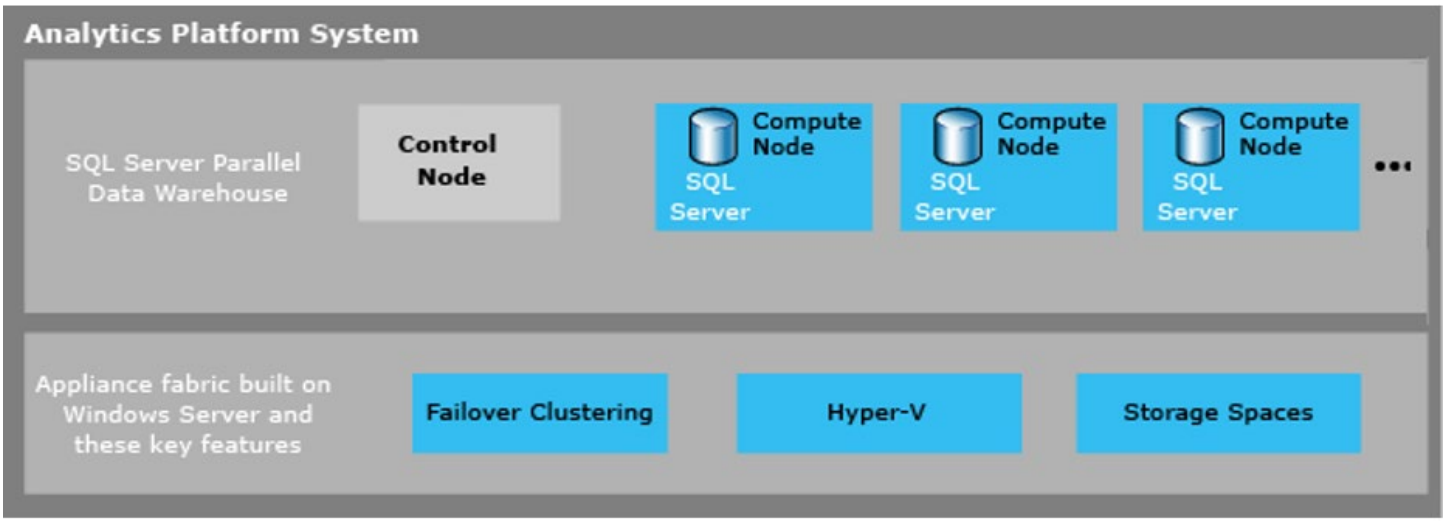

FIGURE 2. Microsoft Analytics Platform System

The details and descriptions provided in this guide, pertain to the current CU7.7 release. New features and enhancements provided with the CU7.7 release are described in this section.

Changes for the Microsoft Analytics Platform System (APS), Cumulative Update 7 (CU7.7) release

Overview

With the APS CU7.7 release, the HPE ConvergedSystem 300 for Microsoft Analytics Platform adds a number of new features and modifications to existing ones. Below is a summary of these changes:

Customer features

- APS CU7.7 software upgrades VMM VM to Windows Server 2016 and installs SCVMM2016. SCVMM 2012 R2 that is currently in use has an end of life date of July 2022. The newer SCVMM is needed to be supported making CU7.7 a mandatory upgrade. Customers are urged to upgrade to CU7.7 as soon as possible.
- New APS SSIS destination adapter that supports SQL Server 2019 as deployment target can be downloaded from download site.
- All CU7.1 to CU7.6 features are rolled up into CU7.7. See Microsoft online docs for more information.

UPGRADE ROADMAP

The upgrade paths (Figure 3) to the APS CU7.7 release can vary depending on what release the HPE ConvergedSystem 300 is currently running on. The current upgrade matrix for APS CU7.7 is detailed below. This matrix is valid for Version 2 (V2) versions of the appliance (AU0.5 or greater) only.

NOTE

Version 1 (V1) of the appliance was the HP Enterprise Data Warehouse Appliance. Version 2 (V2) of the appliance was the HP AppSystem for Microsoft SQL Server 2012 Parallel Data Warehouse, and the HPE ConvergedSystem 300 for the Microsoft Analytics Platform (current offering).

NOTE

In some cases, HDI can be upgraded, but only after a full erase and restore of the region. HDI regions are not supported with CU7.7. Customers with existing HDI regions cannot upgrade to CU7.7 unless the HDI region is removed.

	↓	↓	SW Version to Upgrade To↓	↓	↓	↓	↓	↓	↓	↓
Current SW Version ↓	V2 AU1	V2 AU2	V2 AU3	V2 AU4	V2 AU5	V2 A2016	V2 AU7	V2 CU7.5	V2 CU7.6	V2 CU7.7
V2 AU0.5	PDW Only	PDW Only**	PDW Only	N/A	N/A	N/A	N/A	N/A	N/A	N/A
V2 AU1	N/A	PDW Only**	PDW Only	PDW + HDI	PDW + HDI	PDW	PDW	PDW	PDW	PDW
V2 AU2	N/A	N/A	PDW Only	PDW + HDI	PDW + HDI	PDW	PDW	PDW	PDW	PDW
V2 AU3	N/A	N/A	N/A	PDW + HDI	PDW + HDI	PDW	PDW	PDW	PDW	PDW
V2 AU4	N/A	N/A	N/A	N/A	PDW + HDI	PDW	PDW	PDW	PDW	PDW
V2 AU5	N/A	N/A	N/A	N/A	N/A	PDW	PDW	PDW	PDW	PDW
V2 APS2016	N/A	N/A	N/A	N/A	N/A	N/A	PDW	PDW	PDW	PDW
V2 AU7	N/A	N/A	N/A	N/A	N/A	N/A	N/A	PDW	PDW	PDW
V2 CU7.5	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	PDW	PDW
V2 CU7.6	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	PDW

FIGURE 3. Upgrade Matrix

** HDI can be upgraded via a workaround that erases all data stored on the HDI Region

HPE CONVERGEDSYSTEM300 FOR MICROSOFT ANALYTICS PLATFORM SOLUTION ARCHITECTURE DETAILS

Understanding business data continues to be a challenge at most large enterprise companies. The challenges of attaining insight from their existing data, as well as integrating big data sources of information, continue to challenge IT departments. The volumes of data continue to expand, and the answers needed to achieve new business insight require more and more computing power. Data is available in multiple source systems, including non-traditional data sources such as Hadoop data clusters, as well as traditional relational data warehouse systems. Data must be accessed quickly, securely, and reliably to address the needed business insights, ultimately contributing to better insights to the enterprise.

Fortunately, the HPE ConvergedSystem 300 for Microsoft Analytics Platform addresses these needs. The HPE ConvergedSystem 300 for Microsoft Analytics Platform solution provides a combination of HPE hardware, HPE software, and software from Microsoft to provide a

Massively Parallel Processing (MPP) version of the Microsoft SQL Server platform, known as Microsoft SQL Server Parallel Data Warehouse. This solution provides customers the ability to handle very large volumes of data (up to 12.6 Petabytes) in a single solution, providing lightning fast query responses to even the most difficult questions. Fast data loading and secure data storage solutions combine to provide an enterprise ready data warehousing platform. Additionally, the Microsoft Business Intelligence (BI) solutions integrate quickly and easily into the HPE ConvergedSystem 300 for Microsoft Analytics Platform solution, providing an enterprise data warehousing framework for reporting, data analysis, and self-service BI, even allowing Microsoft Excel query integration directly into the HPE ConvergedSystem 300 for Microsoft Analytics Platform.

Part of the infrastructure of many big data solutions has become the rapidly expanding world of Hadoop. SQL Server Parallel Data Warehouse introduced and includes the PolyBase query engine to combine Hadoop data into the existing SQL Server Parallel Data Warehouse relational data store. PolyBase can be used to query external Hadoop data sources.

The HPE ConvergedSystem 300 for Microsoft Analytics Platform solution offering includes the following:

- HPE Infrastructure
  - Two InfiniBand switches for fast data transfers and redundancy
  - Two Ethernet switches for management and querying capabilities to the solution
  - HPE ProLiant DL360 Gen10 servers running Microsoft Analytics Platform System software
  - HPE D6020 disk enclosures to service the very large data volumes
  - Intelligent power modules for redundant power to components
- Microsoft Analytics Platform System
  - Software to provide the SQL Server relational database services in the HPE ConvergedSystem 300 for Microsoft Analytics Platform
  - Data Management software to control parallel operations in the Massively Parallel Processing (MPP) appliance, and Microsoft PolyBase which enables seamless querying of data.
  - Management software, including Active Directory and System Center for internal management of the solution
  - Client tools for secure, reliable data access
- HPE Management Software
  - The Support Pack for the HPE ConvergedSystem 300 for Microsoft Analytics Platform solution provides:
    - Tools for ConvergedSystem validation and diagnostics
    - Reporting tools for health and proactive care
    - Validated Firmware/Driver packages for version control and updates
  - HPE Insight Remote Support for proactive customer care
- HPE Services
  - A collaborative support model with Microsoft, for a single phone call to support, for all aspects of the solution
  - Factory integration and build of all components, including software deployment
  - Delivered fully built and integrated
  - On-site installation services included with the solution

Design objectives

This section outlines key design objectives for the HPE ConvergedSystem 300 for Microsoft Analytics Platform.

Leadership performance

Powered by HPE ProLiant DL360 Gen10 (2 x Intel Xeon-Gold 6134) servers, the solution provides excellent performance and scalability. The solution can be expanded with up to 56 Active Scale Unit nodes or 64 total nodes. All these servers work together to provide Massively Parallel Processing (MPP) through a single, integrated interface.

Realistic sizing

The HPE ConvergedSystem 300 for Microsoft Analytics Platform must start with a Base Scale Unit (1/4 rack) either with 151TB or 453TB of User Data capacity (based on 2TB or 6TB disks, 5:1 compression), scaling up to 4PB or 12.6PB (based on 2TB or 6TB disks, 5:1 compression) of storage on a fully populated solution.

After the initial Base Scale Unit, you can scale up by purchasing additional scale units. This allows a customer to start small and grow their HPE ConvergedSystem 300 for Microsoft Analytics Platform as their data sizing needs grow.

High availability

The HPE ConvergedSystem 300 for Microsoft Analytics Platform solution is designed with high availability at all levels providing maximum uptime. The solution includes:

- Failover clustering: all physical servers are part of a Windows failover cluster, which includes multiple redundant failover servers as the solution scales up
- Failover of the entire infrastructure, which is all contained within virtual machines within the HPE ConvergedSystem 300 for Microsoft Analytics Platform solution
- Mirrored disks for all management servers, host operating systems, and data storage
- Redundant Ethernet switches
- Redundant InfiniBand switches
- Redundant power design with redundant HPE Intelligent Power Distribution Units and power supplies

System components

The HPE ConvergedSystem 300 for Microsoft Analytics Platform delivers the following key components.

Hardware

- HPE 42U 600x1075 G2 Enterprise Shock Rack and needed power and management cables.
- HPE DL360 Gen10 servers with 2x Intel Xeon-Gold 6134 processors, with 256GB memory per server
  - Two HPE 600GB 12G SAS 10K 2.5 inch HDD per server
  - Used for hosting management components and failover capabilities
  - Two servers' minimum
- Up to 28 "Scale Units", each consisting of:
  - Two HPE DL360 Gen10 servers with 2x Intel Xeon-Gold 6134 processors, with 256GB memory per server
    - Two HPE 600GB 12G SAS 10K 2.5 inch HDD per server
  - HPE D6020 disk enclosure with seventy (70) disks per enclosure
    - Disks are either 2TB or 6TB 12G SAS 7.2K 3.5 inch DP MDL HDD

NOTE
All disks in the HPE ConvergedSystem 300 for Microsoft Analytics Platform must be the same size (e.g. either 2 or 6TB). Refer to your Account Manager or Reseller for more details).

- HPE 5900AF48G 4XG 2QSFP+ Switch – Ethernet (2)
- Mellanox InfiniBand EDR v2 100 Gb/sec 36-port Power-side-inlet Airflow Managed Switch (2)
- Choice of HPE Power Distribution Units (PDUs) – Choices include three phase (40A NA/Japan or 32A International) (two per rack) or single phase (40A NA/Japan or 32A International) (four per rack) Modular Power Distribution Units (PDUs); priced separately

Management software
- Windows Server 2012 R2 Standard Edition on all physical servers and virtual machines
- Active Directory for administrative control and security
- Microsoft System Center 2012 Virtual Machine Manager
- Microsoft Analytics Platform System
- HPE Integrated Lights-Out (iLO) Advanced
- HPE Support Pack for the HPE ConvergedSystem 300 for Microsoft Analytics Platform

Services
- HPE Factory Express Deployment Service
- HPE Proactive Care, 3 years included. Additional support offerings are detailed in the QuickSpecs

Base components: Quarter Rack

A "Quarter Rack" (1/4 rack) configuration (Figure 4) is the minimum HPE ConvergedSystem 300 for Microsoft Analytics Platform solution that can be purchased. The solution includes the InfiniBand and Ethernet switches, an orchestration server and a failover server (HST01 and HST02, respectively), and a single scale unit. This configuration is also referred to as a 2 + 2 configuration, because there are two scale unit servers + two failover servers (including the PDW Orchestration node (HST01)).

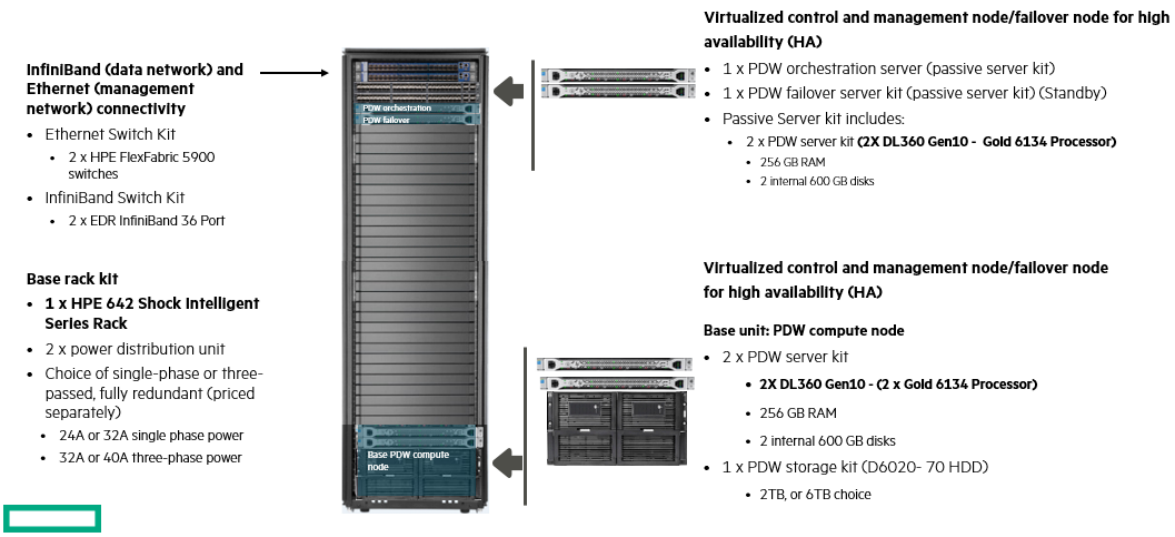

FIGURE 4. Components of a "Base" or "Quarter Rack" system

For a detailed listing of the components included see Appendix A – Configuration details.

This Base Unit is the smallest HPE ConvergedSystem 300 for Microsoft Analytics Platform configuration available, including the minimum amount of hardware and software, and includes the following:

- Two HPE 5900AF48G 4XG 2QSFP+ Switch - Ethernet
- Two Mellanox InfiniBand EDR v2 100 Gb/sec 36-port Power-side-inlet Airflow Managed Switch
- Two management servers, HST01 (the primary Orchestration Server) and HST02 (the primary Failover Server)
- Two Scale Unit servers, HSA01 and HSA02, along with one D6020 disk enclosure

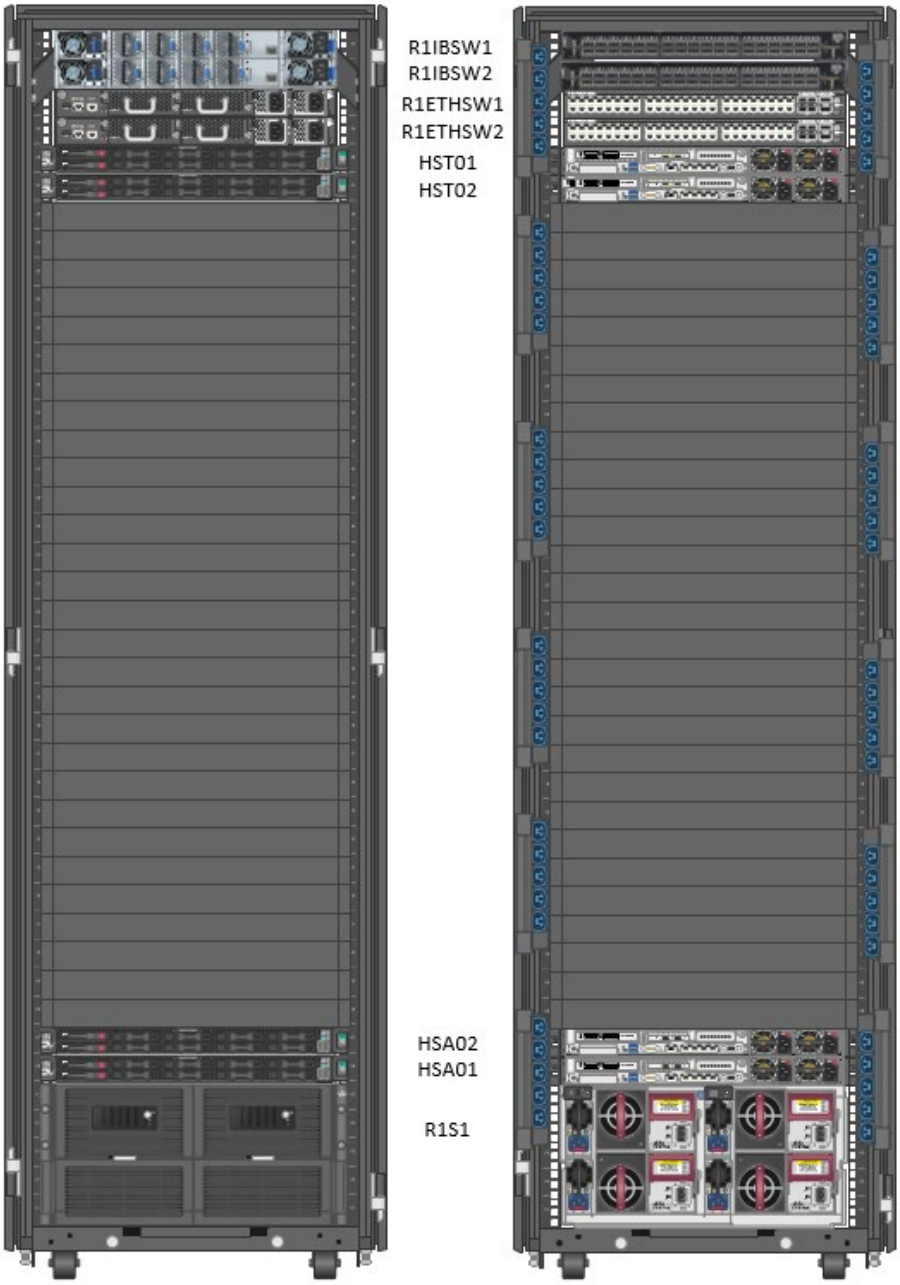

FIGURE 5. Front and rear views of a "Quarter Rack" HPE ConvergedSystem 300 for Microsoft Analytics Platform

The systems in this base scale unit perform the following functions:

- HST01 runs three virtual machines and one service:
  - CTL01 – "Control Node", runs the Massively Parallel Processing (MPP) engine, controls Data Movement Service (DMS) on all nodes, and is the location where client applications connect.
  - AD – Active Directory for the Fabric Domain, owning all hardware and iSCSI virtual machines.
  - VMM – Hosts System Center 2012 Virtual Machine Manager, performing VM provisioning and management. Also provides Windows Server® Update Services (WSUS) for Windows Update.
- HST02 does not actively run a workload, the server is available to host one or more virtual machines in the event of some kind of failure in the environment.
- HSA01 runs CMP01, the PDW "compute node" virtual machine that runs the SQL Server PDW workload. It also hosts the iSCSI01 VM to coordinate storage within the solution.
- HSA02 runs CMP02, another PDW "compute node", as well as the iSCSI02 VM.
- The D6020 disk enclosure contains seventy (70) 2TB or 6TB drives. These drives hold the PDW relational data, as well as the virtual machines hosted on HST01, and all the data accessed by HSA01 and HSA02 virtual machines.

Large "Full Rack" solution

A fully populated HPE ConvergedSystem 300 for Microsoft Analytics Platform configuration will contain all the systems in the Base Scale Unit, and may include the following optional components:

- An additional PDW Failover Server. This server will be known as HST03. This server will be physically racked together (shown in Figure 6 below) with the other HST servers.
- Up to three additional scale units, with each scale unit made up of two (2) HPE ProLiant DL360 Gen10 servers, (HSAxx), along with one D6020 disk enclosure.

The additional systems in this configuration perform the following functions:

- HST03 does not actively run a workload; the server is available to host one or more virtual machines in the event of some kind of failure in the environment of HST01 or 02.
- HSAxx servers will run the PDW compute node virtual machines, named CMPxx. All HSAxx servers run an iSCSI virtual machine as well.
- Each D6020 disk enclosure contains seventy (70) 2TB or 6TB drives. These drives hold the PDW relational data for each set of virtual machines hosted on the attached HSAxx servers.

In Figure 6 below, a fully populated rack is shown. This system has the optional HST03 server as a secondary failover server for the PDW workload. Four Scale Units are installed, the maximum capacity for a single rack.

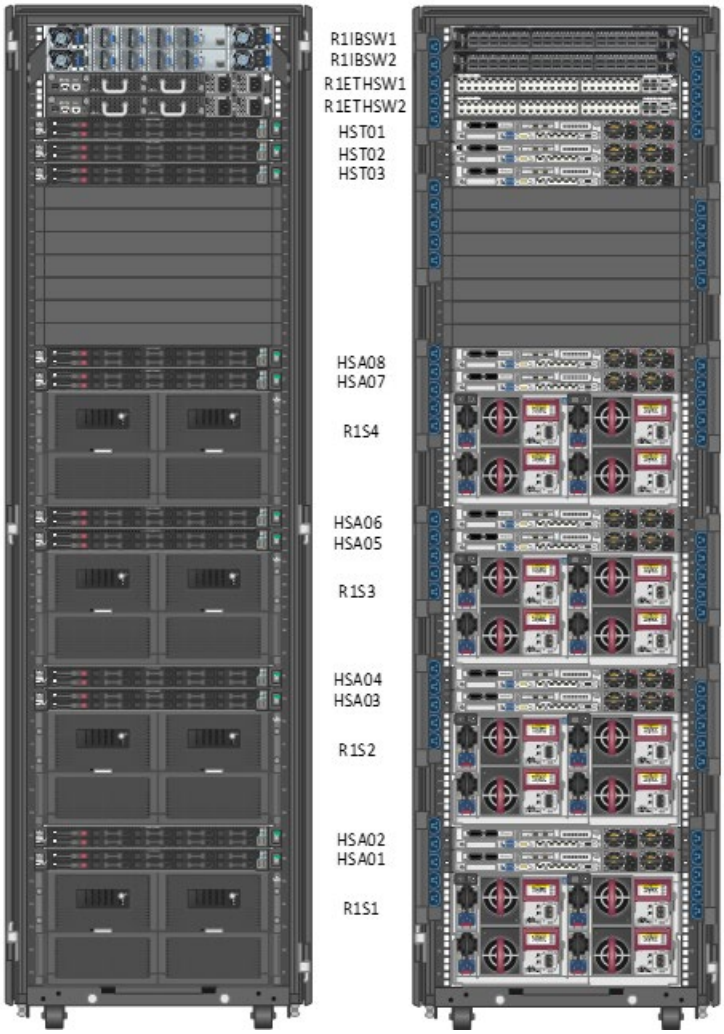

FIGURE 6. Front and rear views of a "Full Rack" in an HPE ConvergedSystem 300 for Microsoft Analytics Platform

In this "Full Rack" configuration (Figure 6), all possible servers for an HPE ConvergedSystem 300 for Microsoft Analytics Platform solution that can be purchased for a rack are included (solution can scale up to seven racks). The solution includes the InfiniBand and Ethernet switches, a management server (HST01), two failover servers (HST02 and HST03), and four scale units.

This configuration is also referred to as an 8+3 configuration, because there are eight scale unit servers + three failover servers.

NOTE

If this were a secondary rack, the scale units would be identical, along with the InfiniBand and Ethernet switches at the top. However, there would be at most a single failover server at the top of the rack, in the position of HST01 in this diagram. In the secondary (and other additional) racks, however, the server name would simply increment. In the Full Rack example above, when a second rack of hardware is added to this solution, the failover server at the top of the rack would be HST04 (since that's the next available number in the sequence).

The HSAxx servers would continue to be counted from the bottom of the rack, HSA09-HSA16 in rack #2. This pattern continues through the remaining racks of a single HPE ConvergedSystem 300 for Microsoft Analytics Platform solution.

Regions

From a software perspective, the HPE ConvergedSystem 300 for Microsoft Analytics Platform is broken up into two "regions." These regions are the PDW region, running the SQL Server Parallel Data Warehouse software, and the Fabric region, providing the management infrastructure.

The PDW region

The PDW region is the set of virtual machines that entail the Microsoft SQL Server Parallel Data Warehouse solution. The PDW region is the Massively Parallel Processing (MPP) software that provides a relational data warehouse for the customer. The Orchestration Server (HST01), or control or management node is the primary access/query point for customers. The Orchestration Server coordinates querying and responses from the compute nodes (CMPxx virtual machines) running on the HPE ConvergedSystem 300 for Microsoft Analytics Platform scale unit servers (Figure 7), running one virtual compute node per physical server in the scale unit.

The orchestration node additionally runs some of the Fabric region VMs such as AD, CTL, VMM, and WDS.

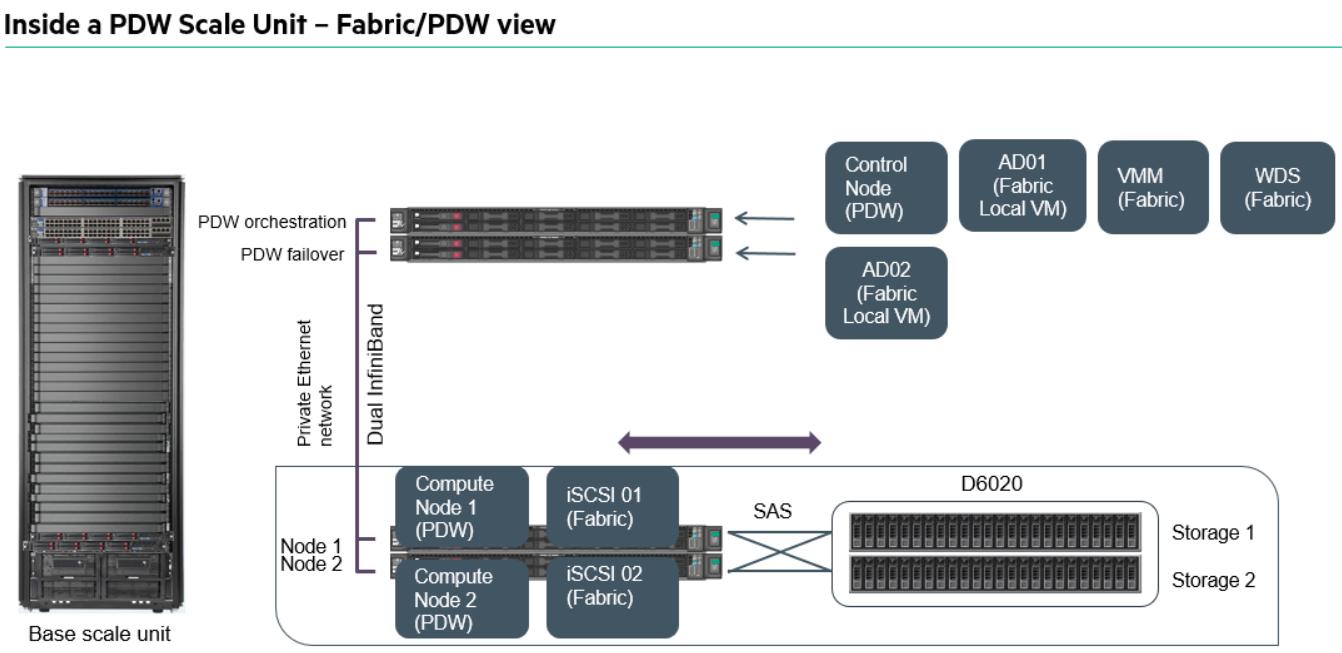

FIGURE 7. Orchestration server review

Support services are run on the System Center 2012 Virtual Machine Manager VM (VMM virtual machine), providing multiple layers of management support to the virtual and physical infrastructure.

The PDW Failover Server (HST02) can also run these components in a failover scenario, as well as running a compute node (CMPxx) virtual machine in the event of a failure of an HSA component. These, as noted earlier, will always be the first two servers near the top of the rack on the first and/or only rack.

An optional third failover server can be added to the configuration. Both HST02 and the third failover server (HST03), will normally be strictly in standby mode and not running any clustered virtual machines.

Software stack

Several software components are included with the HPE ConvergedSystem300 for Microsoft Analytics Platform appliance. These software components provide the basic solution framework and infrastructure for the appliance. Additionally, several software components are used for

loading and querying the appliance and are not included in the base software stack. Figure 8 below details the PDW region software components along with some of the components that are used to access the appliance.

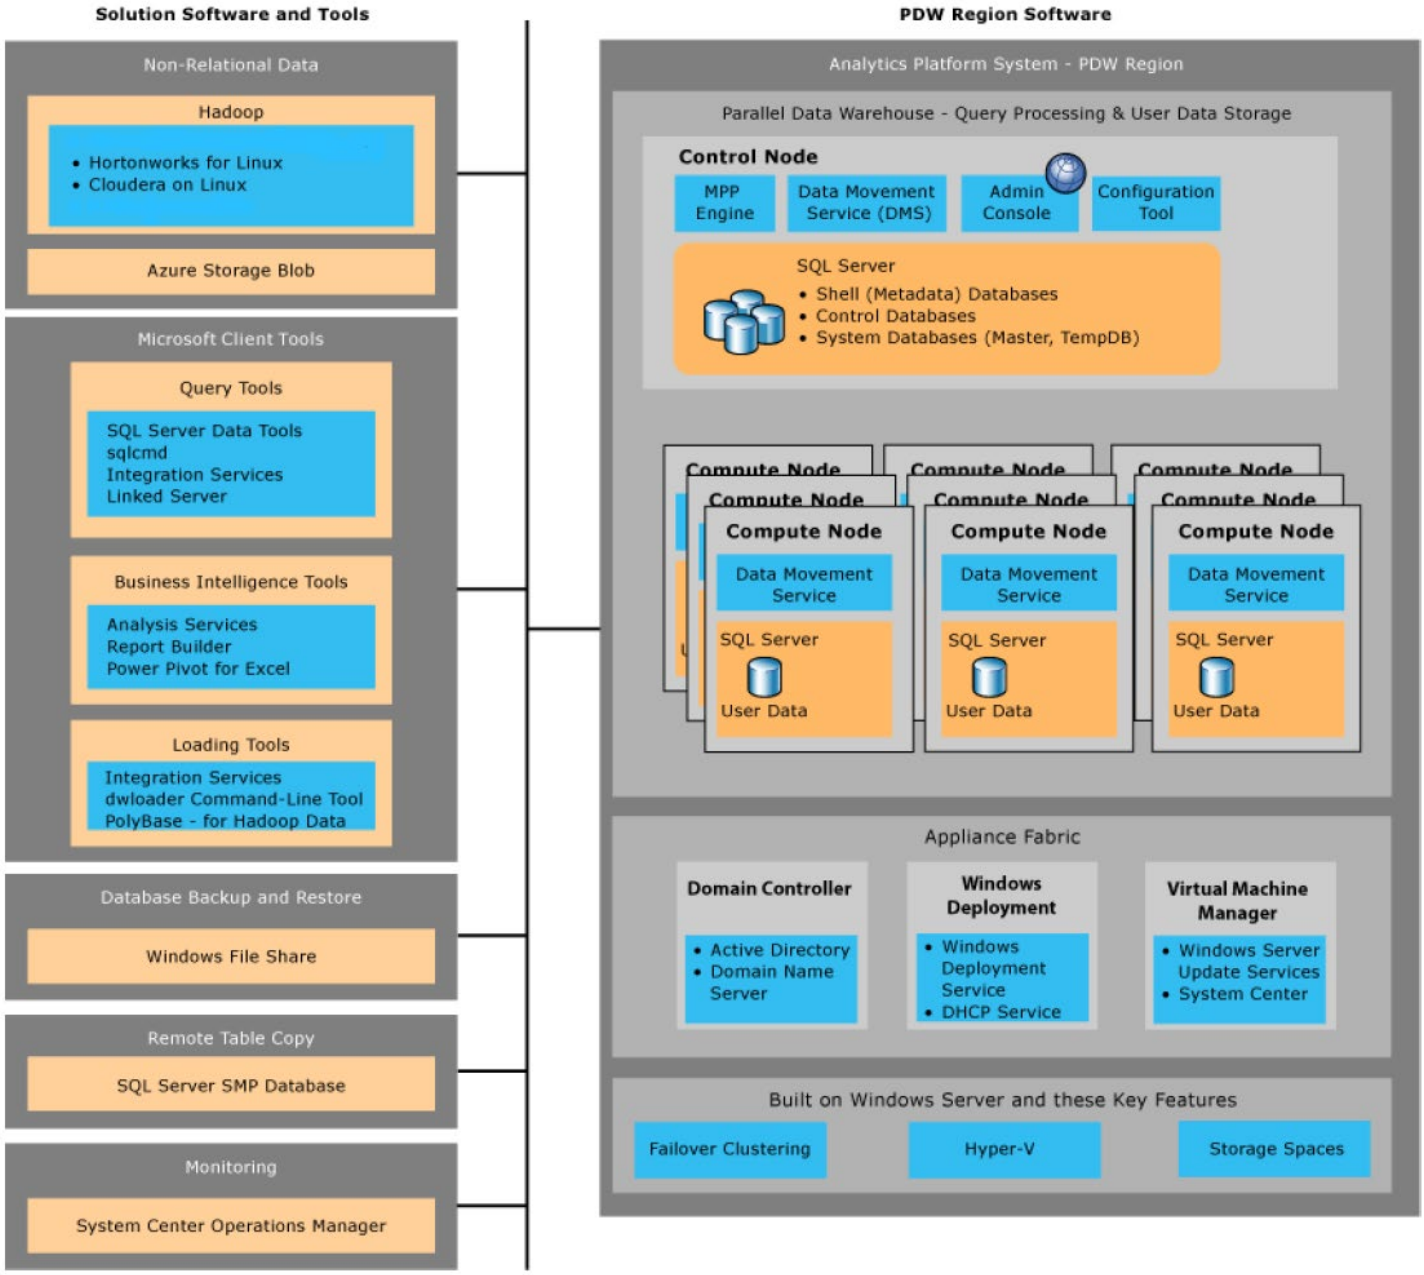

FIGURE 8. HPE ConvergedSystem 300 for Microsoft Analytics Platform (appliance) included software and off-appliance software

Hardware fault tolerance

The HPE ConvergedSystem 300 for Microsoft Analytics Platform uses redundant HPE ProLiant DL360 Gen10 servers as management servers, delivering exemplary performance and efficiency in a compact 1U size. This server is designed for high-performance workloads; with the adaptability you need to satisfy future requirements. In HPE ConvergedSystem 300 for Microsoft Analytics Platform, each HPE ProLiant DL360 Gen10 server has two Intel Xeon-Gold 6134 Processor (8-core 3.2 GHz processors with 25MB cache) and 256GB RAM.

Each server has redundant connections to both the Ethernet and InfiniBand networks, as well as redundant power supplies. The local storage for the operating system consists of two HPE 600GB SAS 10K RPM 2.5 inch HDD in a mirrored (RAID 1) configuration for additional protection.

Scalability

Figure 9 shows an example of a multi-rack HPE ConvergedSystem 300 for Microsoft Analytics Platform. Note that after the first rack, there will be at most one additional failover server per rack, matching the location of HST01. The name of this server will be the next sequentially available HST numbered system, starting at rack one and counting up as you add racks.

Each additional rack can host up to 4 additional scale units and will always have two InfiniBand switches and two Ethernet switches.

NOTE

Environments vary and can contain different quantities of scale units and failover servers. HPE ConvergedSystem 300 for Microsoft Analytics Platform scales up to seven racks containing a maximum of 56 scale units and 64 active and failover server nodes.

FIGURE 9. The HPE ConvergedSystem 300 for Microsoft Analytics Platform solution scales from a "Quarter Rack" up to seven racks

HPE racks and configures HPE ConvergedSystem 300 for Microsoft Analytics Platform in the factory, and installs all the supplied software, including Windows Server 2012 R2 as the bare metal host operating systems, and all virtual machines. This configuration will include the Active Directory configuration for the assigned domains, and the IP address configurations as requested by the customer.

All D6020s and disk sizes within the D6020s must be the same. If they are different, you can only use the amount of space available on the smallest hard disk used in any of the D6020 enclosures.

HPE ConvergedSystem 300 for Microsoft Analytics Platform architecture enables simultaneous query execution, and each individual query can execute on all the servers and all the D6020 disks storing user data, in parallel. This Massively Parallel Processing (MPP) architecture gives HPE ConvergedSystem 300 for Microsoft Analytics the speed to provide user queries with significantly higher levels of I/O throughput than is available on a scale-up SMP system.

For more information, see Architecture.

Data scale units

A data scale unit, as noted above, consists of two HPE ProLiant DL360 Gen10 servers and an HPE D6020 disk enclosure, and is the basic unit of storage and processing scalability.

Data scale unit servers

The HPE ConvergedSystem 300 for Microsoft Analytics Platform uses two DL360 Gen10 servers as scale unit servers, delivering exemplary performance and efficiency in a compact 1U size. This server is designed for high-performance workloads; with the adaptability you need to satisfy future requirements. Each server runs a separate, virtualized workload, and provide connectivity to the HPE D6020 disk enclosure.

Data scale unit storage

The HPE D6020 disk enclosure provides data storage within each scale unit. The D6020 is populated with seventy (70) HPE 12G SAS 7.2K RPM DP MDL HDD, in 2TB or 6TB size to provide each HPE D6020 disk enclosure with 151TB or 453TB of user data capacity respectively.

NOTE
User data capacity is calculated based on disks allocated to database space, factored in the formatted capacity and RAID1, as well as 5x compression ratio.

NOTE
SQL Server Parallel Data Warehouse uses the newest version of in-memory columnstore, which is both updateable and clustered. By reorienting data in a column rather than a traditional row store, organizations can see next-generation performance (up to 50 times) by reducing the query time from hours to seconds and high compression on their data (up to 15 times). This is due to the grouping of data in a columnar format (containing values from multiple rows) that allows for higher efficiency in both storage and the returned dataset. The ability to permanently convert traditional row store tables with the updateable columnstore of xVelocity is another feature, which dramatically reduces storage (by up to 70 percent) while still providing real-time data warehouse query performance.

With the 12Gb/s SAS HPE D6020 disk enclosure, Hewlett Packard Enterprise extends and redefines direct-attached storage, combining the simplicity and cost effectiveness of direct-attached storage without sacrificing flexibility or performance. The HPE D6020 is used as part of a straight-forward in-rack 12Gb/s SAS implementation that delivers high-density, low-cost external zoned direct attach storage for HPE DL360 Gen10 servers.

The HPE ConvergedSystem 300 for Microsoft Analytics Platform solution uses Windows Server Storage Spaces to provide management of the drives in the HPE D6020 disk enclosure.

User data capacity

The HPE ConvergedSystem 300 for Microsoft Analytics Platform can scale up to 4 or 12.6 Petabytes based on either the 2TB or 6TB D6020 drive option. Figure 10 below details the storage capacity sizes using different disk options.

User data capacity	User data (TB) – Single rack			
Topology	2 active servers	4 active servers	6 active servers	8 active servers
2 TB disks	151 TB	302 TB	453 TB	605 TB
6 TB disks	453 TB	906 TB	1359 TB	1815 TB

User data capacity	User Data (TB) – Base rack + expansion rack(s)					
Topology	16 active servers (2 racks)	24 active servers (3 racks)	32 active servers (4 racks)	40 active servers (5 racks)	48 active servers (6 racks)	56 active servers (7 racks)
2 TB disks	1209 TB	1814 TB	2418 TB	3023 TB	3627 TB	4232 TB
6 TB disks	3627 TB	5442 TB	7254 TB	9069 TB	10881 TB	12696 TB

Note: User data capacity is calculated based on disks allocated to database space, factored in the formatted capacity and RAID1, as well as a 5X compression ratio.

FIGURE 10. User data capacity based on either 2TB or 6TB disks

NETWORK ARCHITECTURE

Overview

As shown in Figure 11, the HPE ConvergedSystem 300 for Microsoft Analytics Platform includes three logical networks: dual redundant InfiniBand networks for data movement, dual redundant 10GbE networks for customer data, and a unified management network for management traffic and query requests. The networks are integrated into the HPE ProLiant DL360 Gen10 servers through the following components.

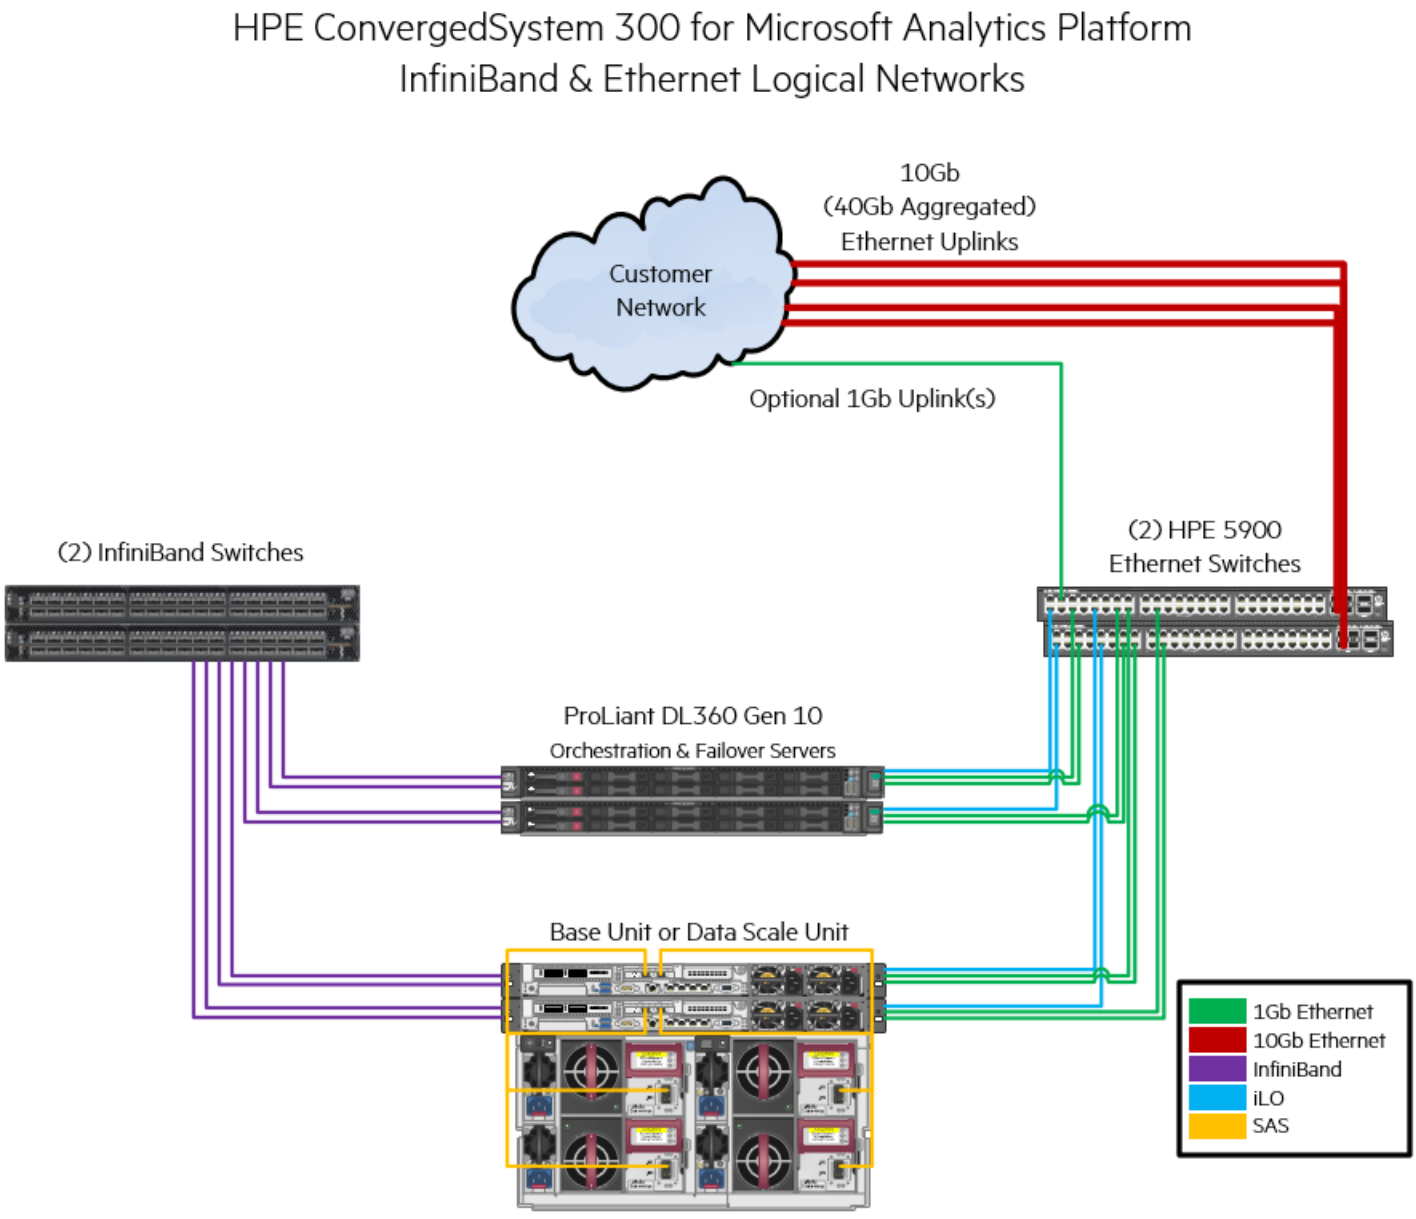

FIGURE 11. Logical InfiniBand and Ethernet Networks

Network components

InfiniBand network

The HPE ConvergedSystem 300 for Microsoft Analytics Platform includes two Mellanox InfiniBand EDR V2 100 Gb/sec 36-Port Power-side-inlet Airflow Managed Switch for each rack. These switches provide high bandwidth data transfers within the solution as necessary for query resolution. Additionally, customers can connect their own servers to the InfiniBand switches to provide high speed data loading and high-speed backup performance to the solution.

All DL360 Gen10 servers in each rack are connected to the InfiniBand switches with the HPE InfiniBand FDR/Ethernet 10Gb/40Gb 2-port 544+QSFP Adapter. The adapter is based on Mellanox ConnectX-3 technology and PCI Express 3.0. They function as a dual ported InfiniBand card for the HPE ConvergedSystem 300 for Microsoft Analytics Platform. Together with the HPE InfiniBand EDR edge switches, they deliver low-latency InfiniBand bandwidth, up to 56Gb/s Fourteen Data Rate (FDR), for performance driven server and storage clustering applications in the HPE ConvergedSystem 300 for Microsoft Analytics Platform.

Every server uses one connection to each InfiniBand switch, using different Ethernet subnets. This provides full redundancy between the two networks in the event of a hardware failure in any of the InfiniBand networking components.

Ethernet network

The HPE ConvergedSystem 300 for Microsoft Analytics Platform includes two HPE 5900AF48G 4XG 2QSFP switches for each rack. These switches provide Ethernet uplinks to the customer network, as well as supporting management traffic, failover clustering support, and general connectivity with each rack and uplinks between multiple racks in the solution. Additionally, customers can connect their own servers to the InfiniBand switches to provide high-speed data loading and high-speed backup performance to the solution.

The HPE 5900AF48G 4XG 2QSFP Switch Series is comprised of Gigabit Ethernet switches that support static Layer 3 routing, diversified services, and IPv6 forwarding, as well as provide four 10-Gigabit Ethernet (10GbE) and two 40-Gigabit Ethernet (40GbE) interfaces. Unique Intelligent Resilient Fabric (IRF) technology creates a virtual fabric by managing several switches as one logical device, which increases network resilience, performance, and availability, while reducing operational complexity. These switches provide Gigabit Ethernet access and can be used at the edge of a network or to connect server clusters in data centers. High availability, simplified management, and comprehensive security control policies are among the key features that distinguish this switch.

All DL360 Gen10 servers in each rack are connected to the Ethernet switches with the HPE Ethernet 1Gb 4-port 331FLR adapter. Each server has three Ethernet connections, including one iLO Ethernet connection, and two redundant Ethernet connections (one to each switch) for general connectivity.

Connectivity

The Ethernet network serves multiple purposes. From a management and interconnect perspective, ports 1-22 of the 1Gb Ethernet connections on the HPE 5900AF48G 4XG 2QSFP switches are used for connectivity to the various servers. Ports 23-47 are available for customer use, including customer network connectivity if the customer desires a 1Gb Ethernet connection to their network.

By default, the HPE ConvergedSystem 300 for Microsoft Analytics Platform solution ships with Intelligent Resilient Framework (IRF) enabled between the two HPE 5900AF48G 4XG 2QSFP switches, so ports 53 and 54 of the 40Gb ports are used for IRF between the two switches. Port 52 of the 10Gb connections can also be used for switch interconnect. The remaining 10Gb ports are available for customer usage. Figure 12 shows the Ethernet switch cable wiring in the first rack of the HPE ConvergedSystem 300 for Microsoft Analytics Platform.

Color key	IB cable lengths		Ethernet cable lengths	
10 GbE Uplink	.5M	2M	1.2M	
Unused	1M	3M	4ft	7ft
User available				
Reserved				

Ethernet Switch #1 (R2ETHSW1) - U40 HPE 5900-48G-4XG-2QSFP+ (JG510A)

1 GbE Ports (1-48)															10 GbE Ports (49-52)		40 GbE Ports (53-54)
1	3	5	7	9	11	13	15	17	19	21	23	25 to 43	45	47	49	51	53
HST03 U36	HST02 U37	HSA01 U6	HSA03 U13	HSA05 U20	HSA07 U27		HST03 iLO U36	HSA04 iLO U14	HSA08 iLO U28								R1ETHSW2 Port 53 U39
2	4	6	8	10	12	14	16	18	20	22	24	26 to 44	46	48	50	52	54
R1PDU1	HST01 U38	HSA02 U7	HSA04 U14	HSA06 U21	HSA08 U28		HST01 iLO U38	HSA02 iLO U7	HSA06 iLO U21	R1IBSW1 U42				R1ETHSW2 Port 48 U39			R1ETHSW2 Port 54 U39

Ethernet Switch #2 (R2ETHSW2) - U39 HPE 5900-48G-4XG-2QSFP+ (JG510A)

1 GbE Ports (1-48)															10 GbE Ports (49-52)		40 GbE Ports (53-54)
1	3	5	7	9	11	13	15	17	19	21	23	25 to 43	45	47	49	51	53
HST03 U36	HST02 U37	HSA01 U6	HSA03 U13	HSA05 U20	HSA07 U27			HSA03 iLO U13	HSA07 iLO U27								R1ETHSW1 Port 53 U40
2	4	6	8	10	12	14	16	18	20	22	24	26 to 44	46	48	50	52	54
R1PDU2	HST01 U38	HSA02 U7	HSA04 U14	HSA06 U21	HSA08 U28		HST02 iLO U37	HSA01 iLO U6	HSA05 iLO U20	R1IBSW2 U41				R1ETHSW1 Port 48 U40			R1ETHSW2 Port 54 U39

FIGURE 12. The Ethernet switches cable wiring in an HPE ConvergedSystem 300 for Microsoft Analytics Platform first rack

If the solution includes more than one rack, then the recommended infrastructure will look like Figure 13 below. The customer connections (if the customer is using 10Gb uplinks to their network) will remain in the first rack and the additional Ethernet switches in the other racks will simply be daisy-chained using ports 50 and 52.

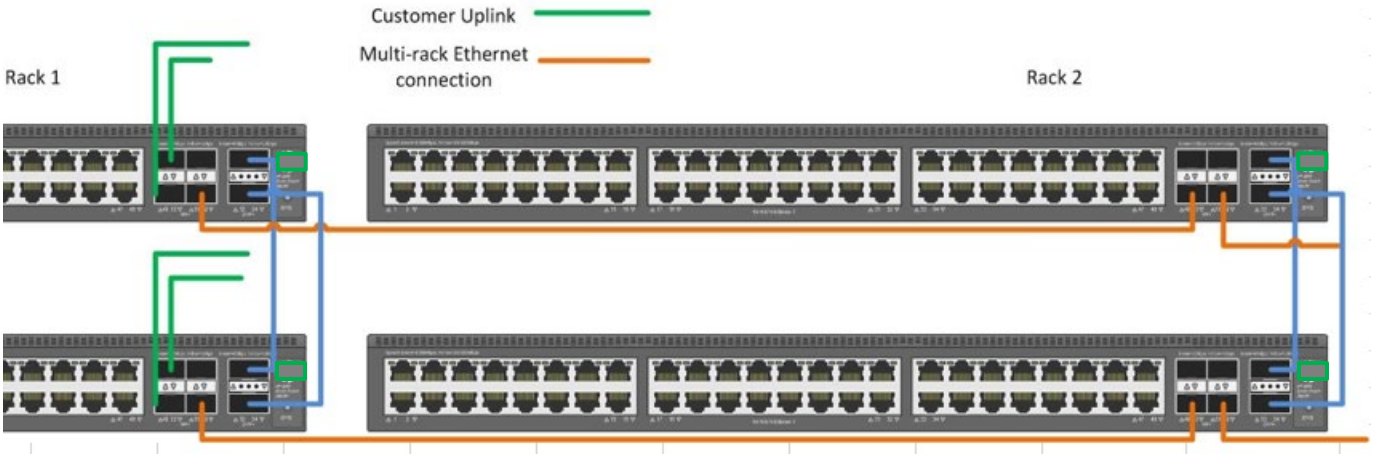

FIGURE 13. Connecting multiple racks of Ethernet switches for the HPE ConvergedSystem 300 for Microsoft Analytics Platform Solution

CLIENT CONNECTIVITY

Client Tools and connection strings

HPE ConvergedSystem 300 for Microsoft Analytics Platform supports Microsoft Analytics Platform Client Tools for the following Data Access APIs:

- ADO.NET
- OLE DB
- ODBC
- JDBC

Supported Query Tools

Microsoft provides and recommends SQL Server Data Tools as the GUI query tool.

Sqlcmd is the SQL Server command-line tool for running T-SQL statements and system commands. It is the recommended command-line tool for querying SQL Server PDW. With sqlcmd you can run T-SQL statements interactively from command-line, as a batch file, or from Windows PowerShell.

Client Tool connectivity – port 17001

ODBC and OLE DB connections use SQL Server Native Client. ADO connections use .NET Framework Data Provider for SQL Server (SqlClient).

All these connections use port 17001, because SQL Server Parallel Data Warehouse only listens to port 17001. To avoid connection failures, append all connection strings with a comma and the port number 17001.

For example, change 10.192.54.48 to 10.192.54.48,17001.

You do not need to rebuild existing SQL Server Integration Services (SSIS) loading packages in order to update the port to 17001.

PolyBase Hadoop integration

PolyBase can query Hadoop data using T-SQL, import data from Hadoop for persistent storage introduced in SQL Server 2016 as a distributed or replicated table, and export data from SQL Server into Hadoop. For more information, see Configure PolyBase Connectivity to External Data.

MANAGEMENT AND MONITORING

HPE ConvergedSystem 300 for Microsoft Analytics Platform provides holistic management for both virtual and physical environments. This section details the different regions of the HPE ConvergedSystem 300 for Microsoft Analytics Platform and the management objects and components for each.

Overview

To understand the management infrastructure, a further understanding of the software infrastructure is necessary.

Admin Console

Overall management is through the Microsoft provided web administration tool within the solution. This web tool is available at the IP address of the cluster of the CTL01 (control node). Within this management interface, overall appliance health, performance monitoring, and region-specific management capabilities are exposed. The Admin Console provides the ability to monitor the following:

- Overall status of the appliance
- Alerts and errors
- Appliance components and their status
- Monitor requests, including queries, loads, backups, and restores
- Monitor additional information for loads, backups, and restores
- Performance information

In Figure 14 below, the home page of the management web page provides a quick summary of the PDW state, along with the option to perform performance monitoring of the PDW region.

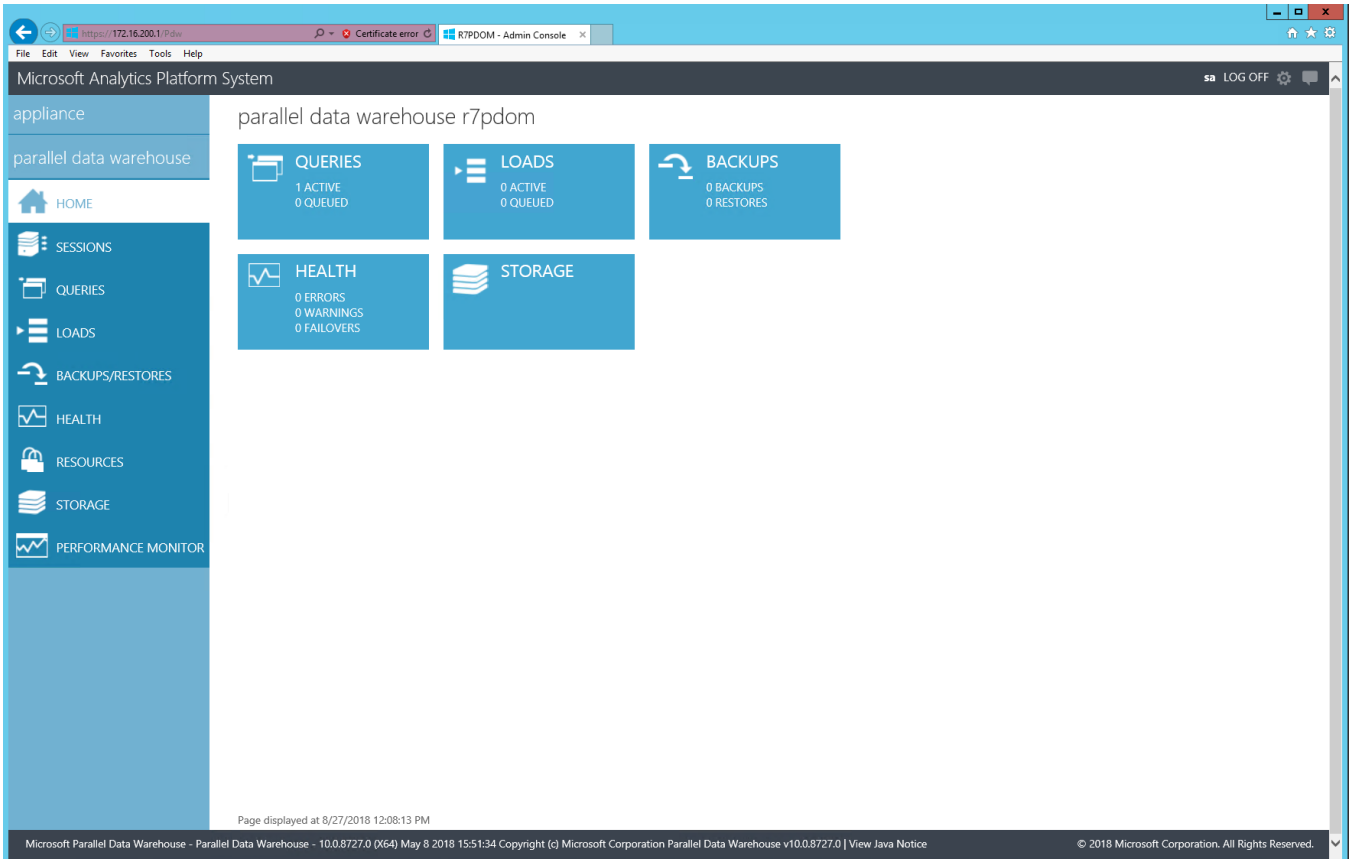

FIGURE 14. The Web Management interface for the HPE ConvergedSystem 300 for Microsoft Analytics Platform

The following sections will cover an overview of the different views of the Admin Console.

Appliance

Home

The Home screen provides a quick summary of the appliance state and is a starting point for viewing the current appliance health and performance.

- Health – Check overall system operation.
- Performance Monitor – Monitor performance in real-time.

Health

The Health screen (Figure 15) provides access to the status of each physical node in the appliance:

- Displays the appliance topology with status indicators showing the health of each monitored component within each node.
- Provides users a detailed view of the current status of individual nodes and properties of the node components.
- Displays hardware and software alerts.
- Displays errors found on each node.

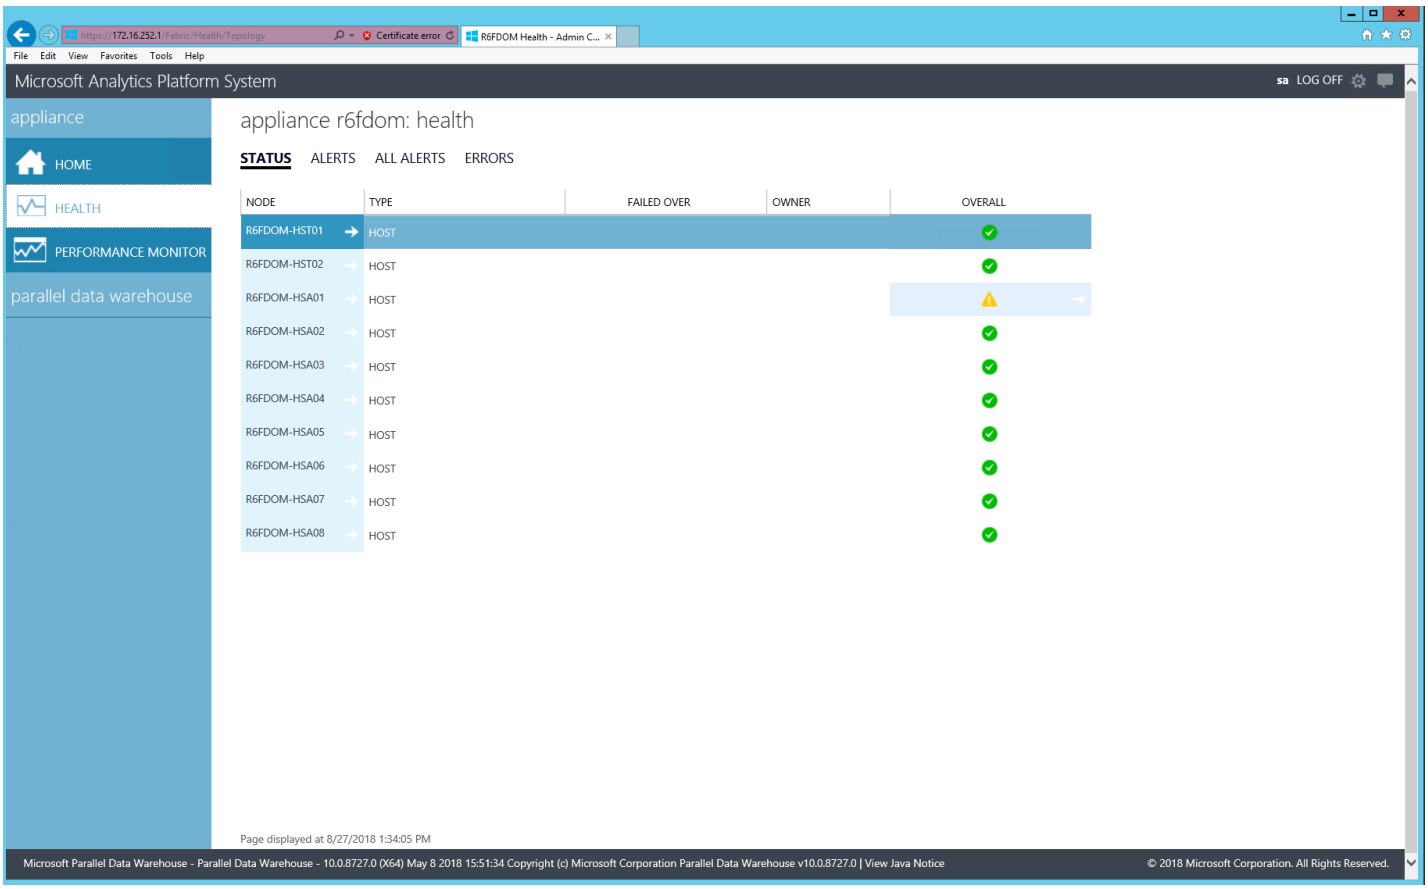

FIGURE 15. Health View for the HPE ConvergedSystem 300 for Microsoft Analytics Platform

Performance Monitor

The Performance Monitor screen (Figure 16) displays the current performance in real time. Various graphs are available for users to monitor various selectable performance counters for logical disks, memory, network interface, process, and processor.

NOTE

Performance Monitor requires the use of Microsoft Silverlight. If your server does not have Silverlight installed, the system prompts you to install it.

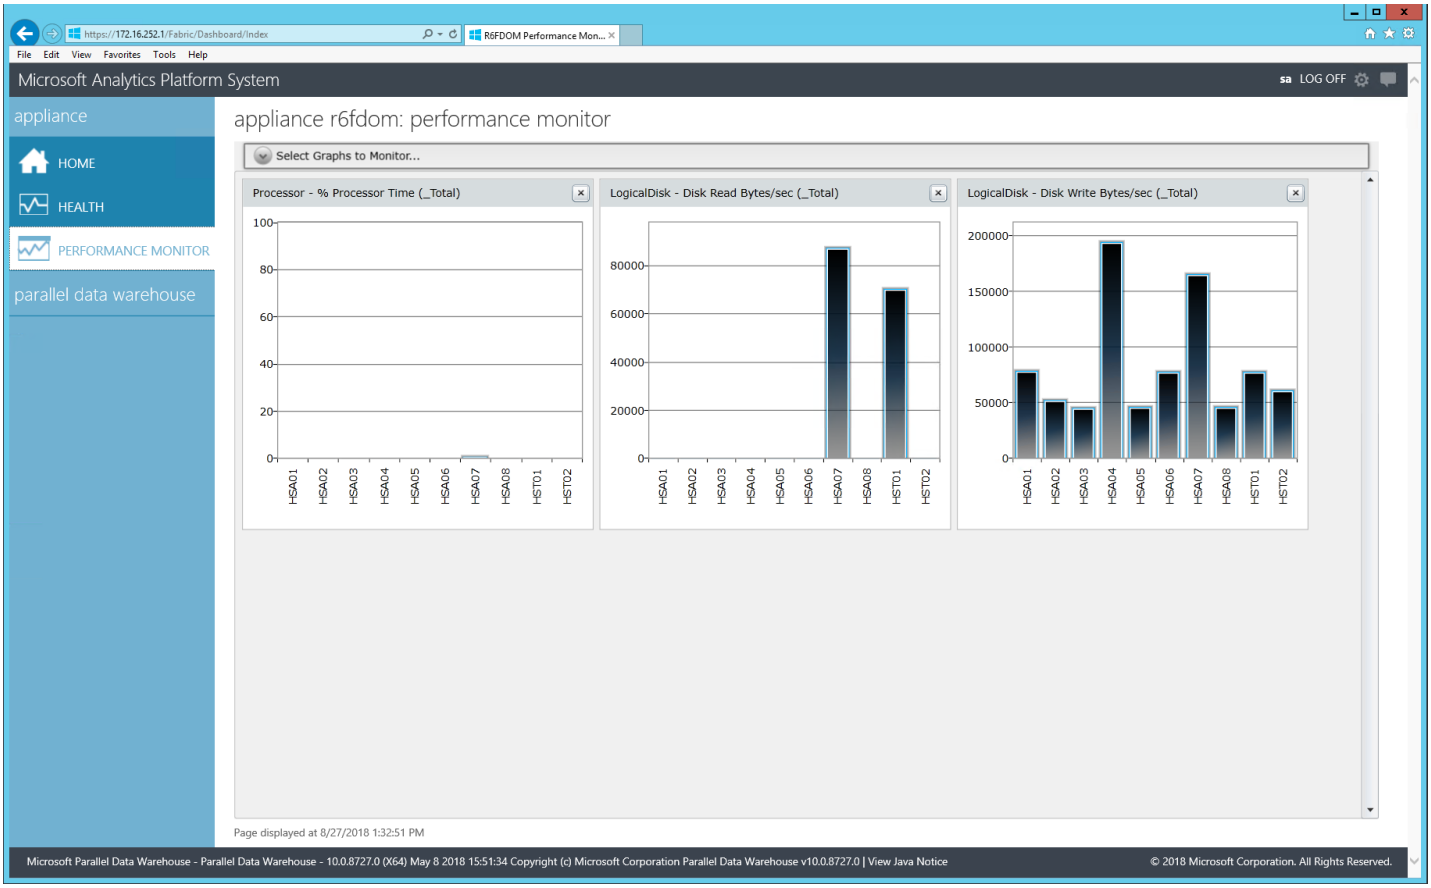

FIGURE 16. Performance Monitor for the HPE ConvergedSystem 300 for Microsoft Analytics Platform

Parallel Data Warehouse (PDW)

Home screen

The Microsoft Analytics Platform - Parallel Data Warehouse Home screen (Figure 17) provides a starting point for viewing and changing PDW configuration information including:

- Sessions – Access individual session information to review activities performed during that session.
- Queries – Review and analyze queries.
- Loads – Check Loads status and detail.
- Backups/Restores – Review current and previous PDW backup and restore operations.
- Health – Check overall PDW Region Health.

- Resources – Check resource status including locks or waits.
- Storage – Review allocation of storage across the PDW Region.
- Performance Monitor – Monitor performance in real-time.

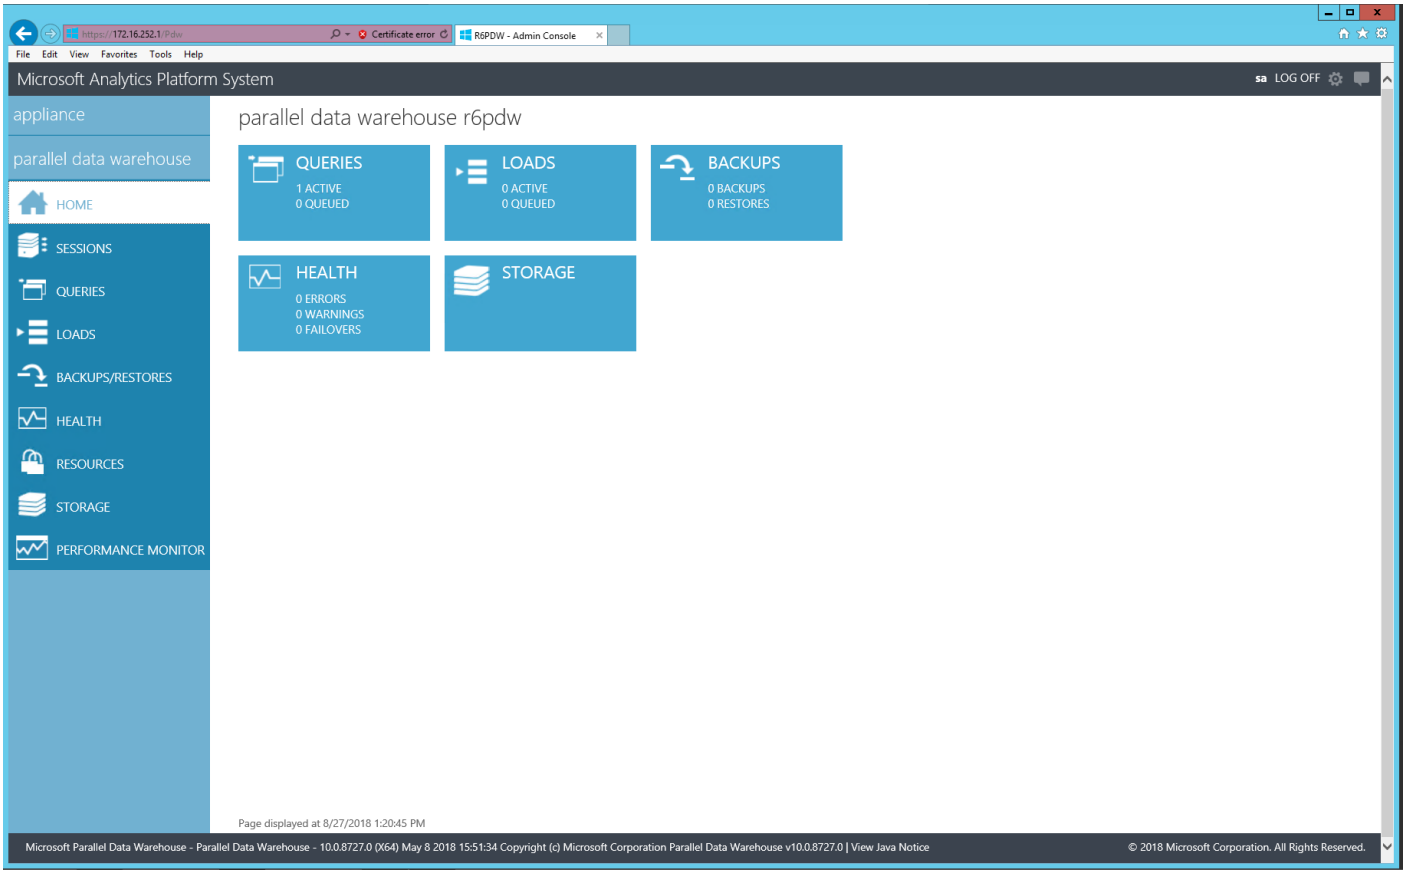

FIGURE 17. Parallel Data Warehouse (PDW) Home screen

Sessions

The Sessions screen (Figure 18) provides a list of user session information. This can help for monitoring resource contention. To view more details, select a SID number in the **SESSION ID** column to display the list of queries, locks, or waits that occurred during that particular session.

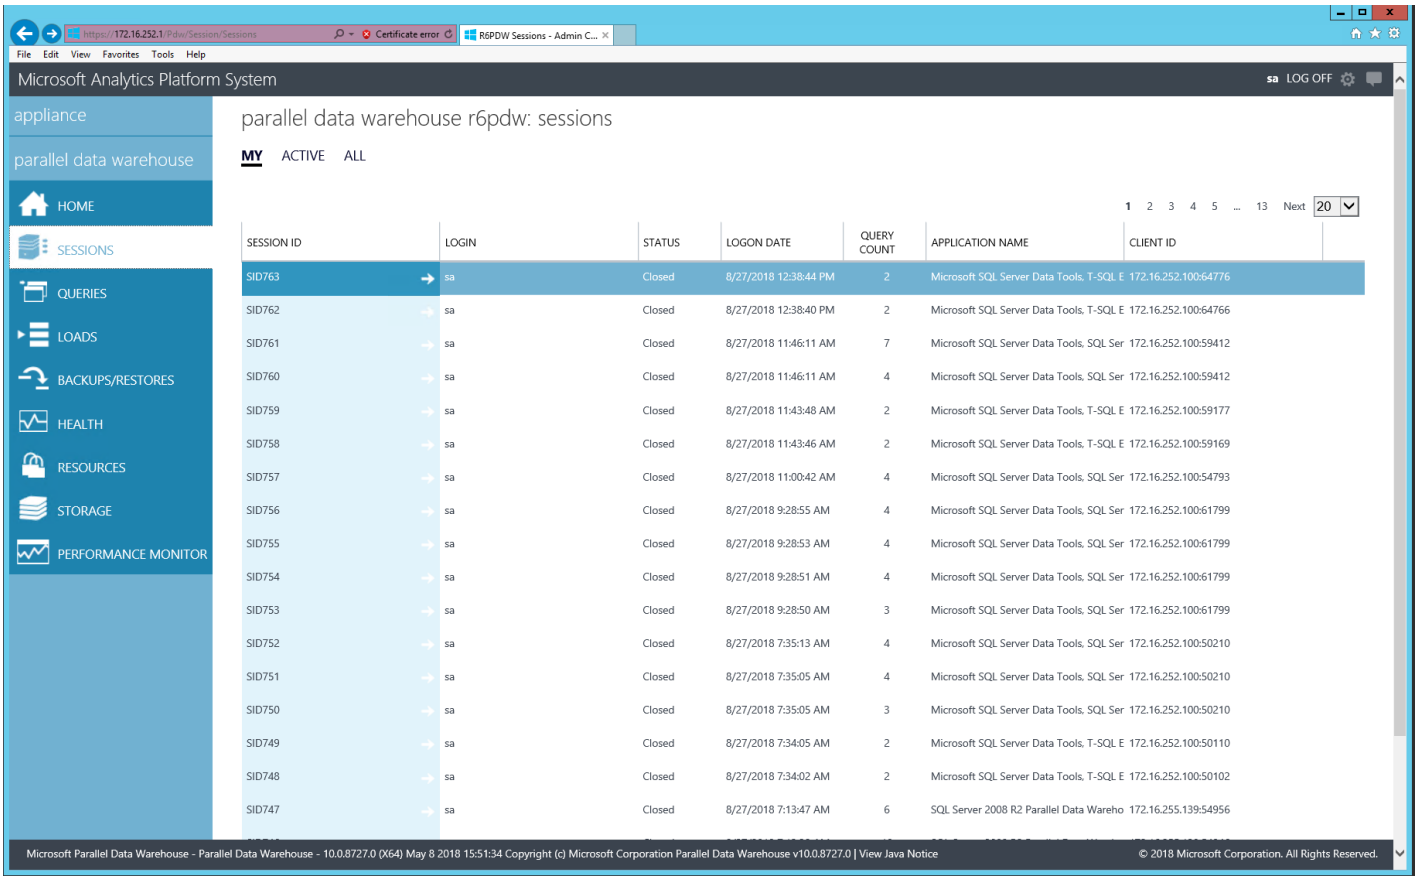

FIGURE 18. HPE ConvergedSystem 300 for Microsoft Analytics Platform – Parallel Data Warehouse (PDW) Sessions

Queries

The Queries screen (Figure 19) provides a list of running queries and recently completed queries. It displays any related errors, as well as providing the details of the query execution plan and the node execution information.

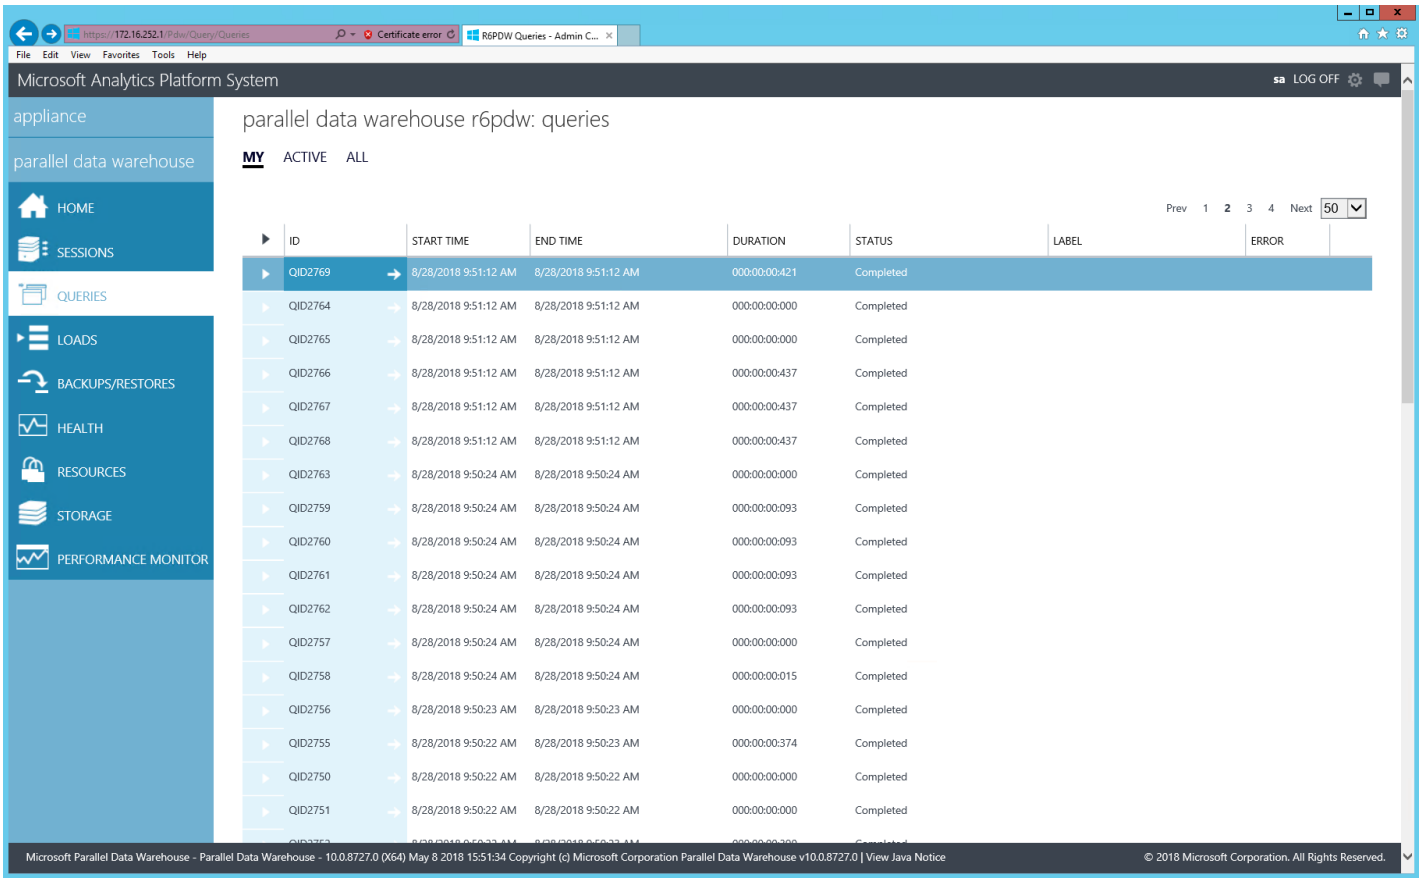

FIGURE 19. Parallel Data Warehouse (PDW) Session Details

For example, selecting query QID2333 (Figure 20), the tool displays information about its query text and executed query plan.

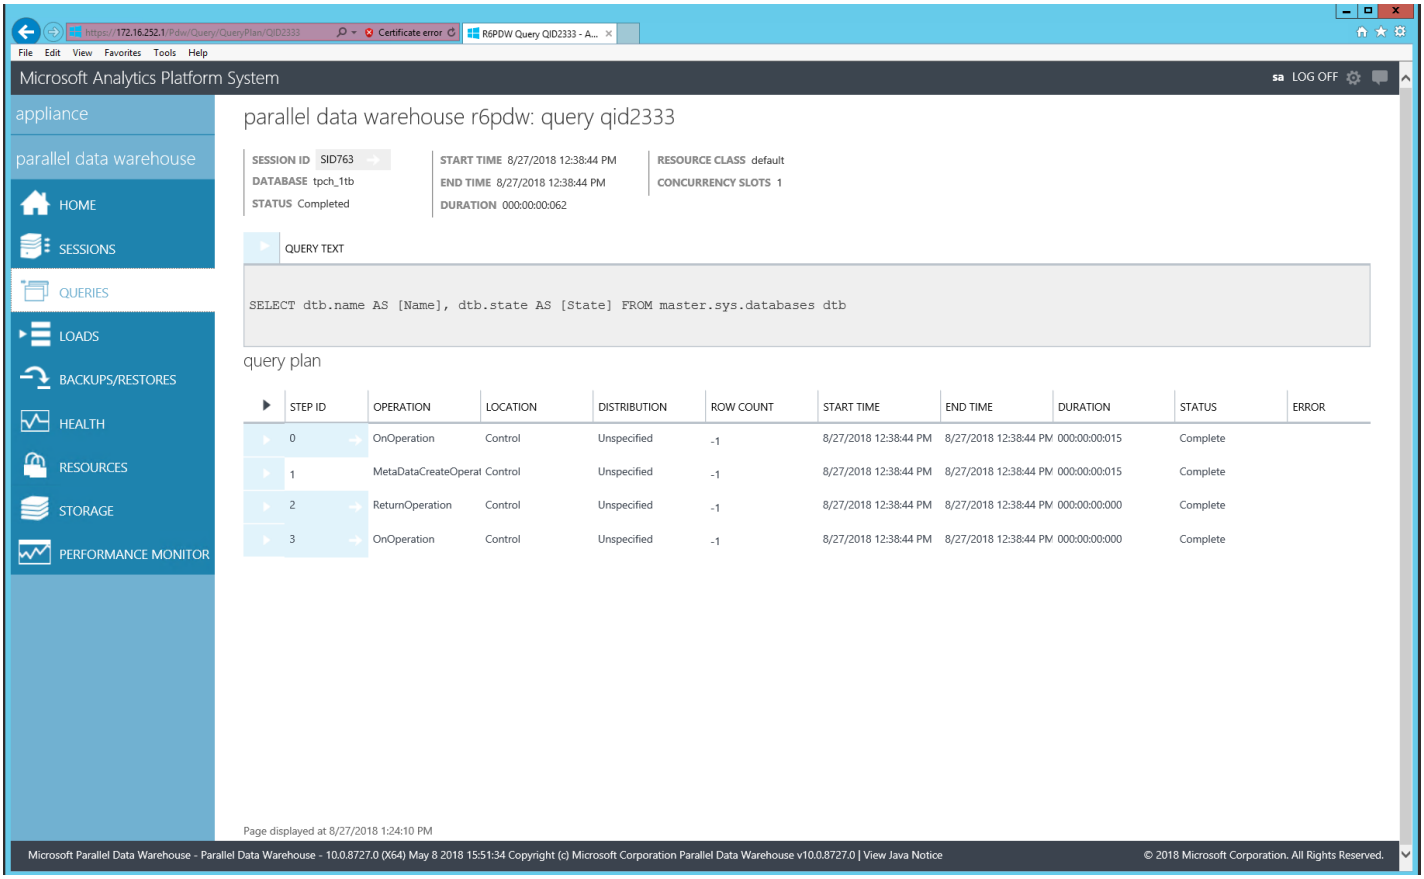

FIGURE 20. Parallel Data Warehouse (PDW) Query Details

You can also view Locks and Waits information from the Resources screen.

Loads

The Loads screen (Figure 21) provides an additional method for accessing Loads information about specific queries. It displays load plans, the current state of PDW loads, and related errors, if any.

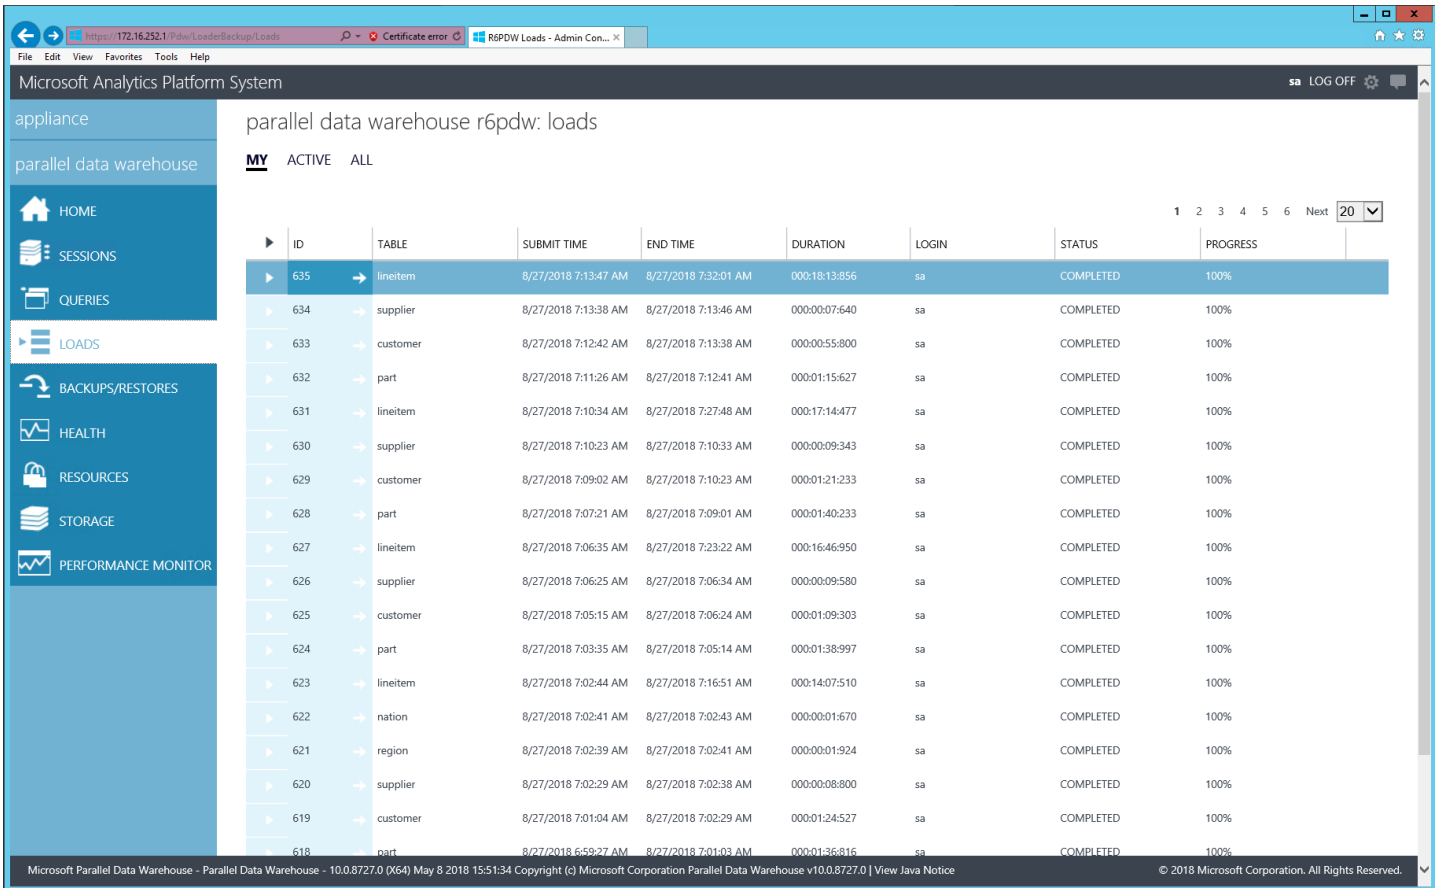

FIGURE 21. Parallel Data Warehouse (PDW) Loads

Selecting an ID number under the ID column (Figure 22) will display the actual query used for that data load, as well as the details associated with the load.

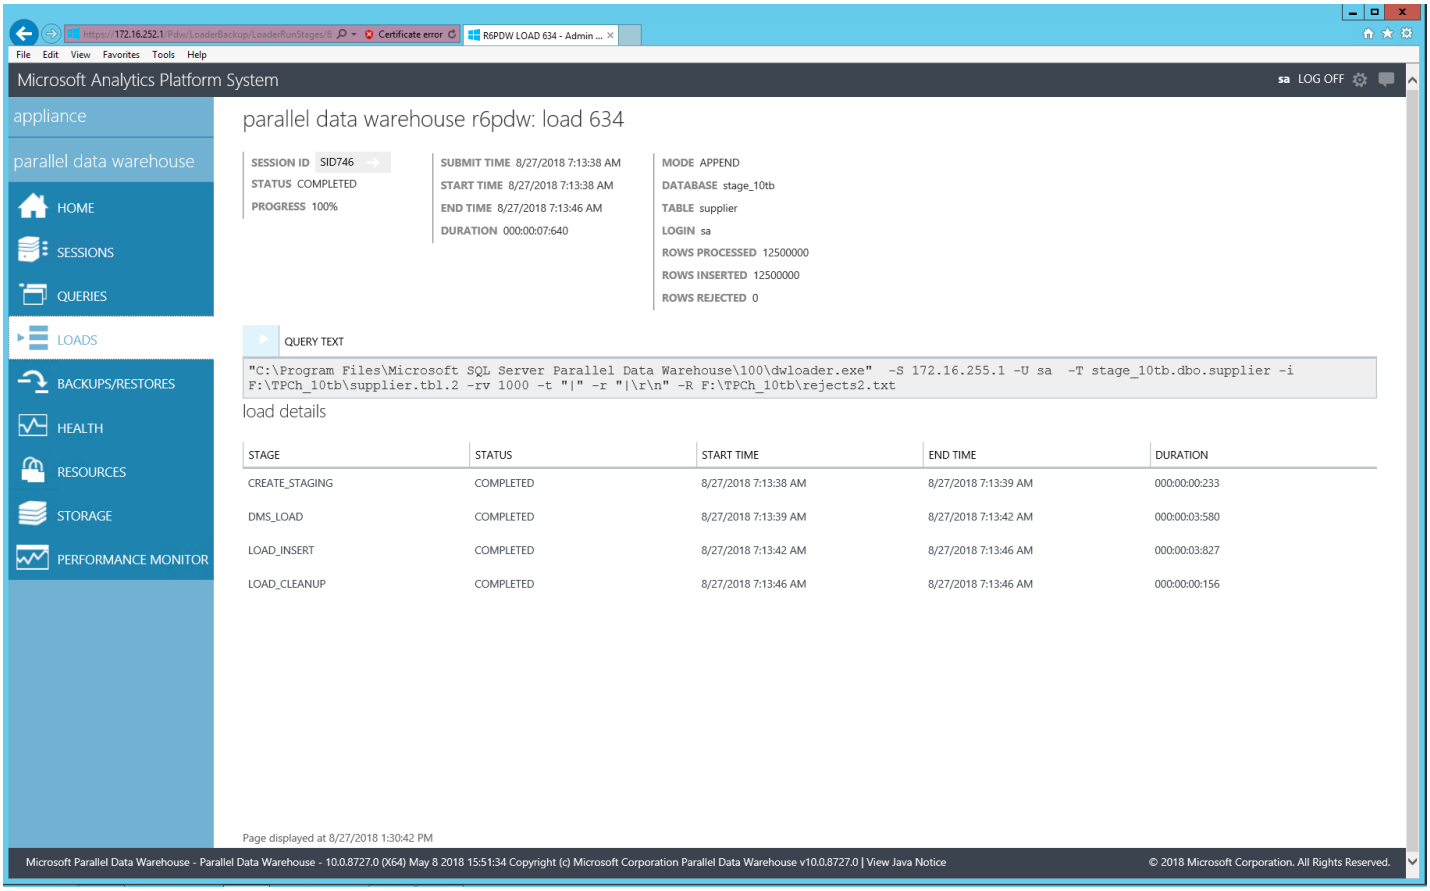

FIGURE 22. Parallel Data Warehouse (PDW) Load Details

Backups/Restores

The Backups/Restores screen (Figure 23) provides a quick view of PDW backup and restore operations. The user interface shows both backup or restore activity per the "Type" column.

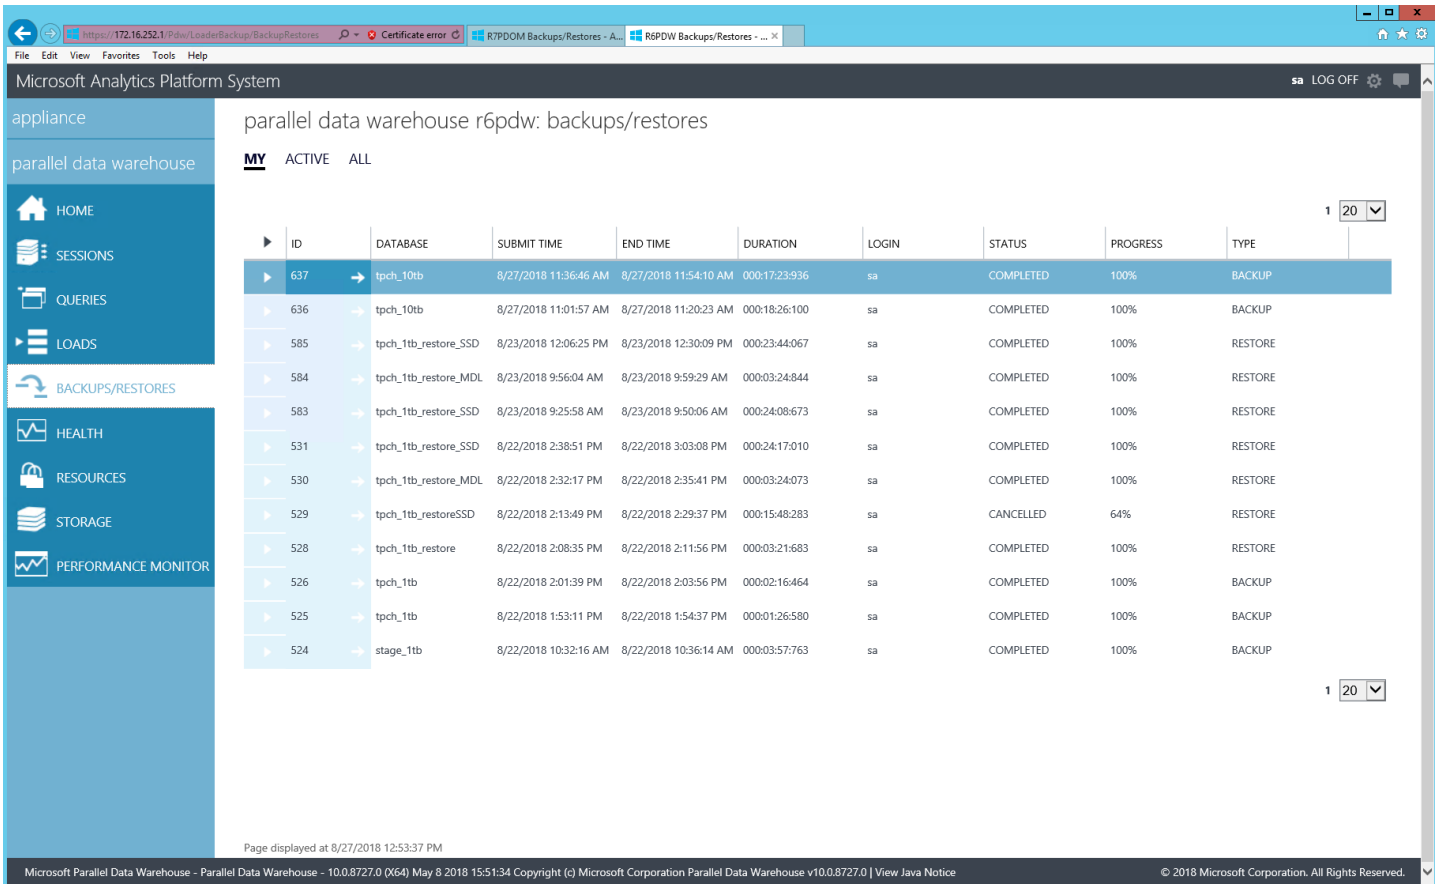

FIGURE 23. Parallel Data Warehouse (PDW) Backups/Restores

Selecting the ID will return the details of the backup or restore activity. When you issue a backup or restore command, the SQL Server Parallel Data Warehouse software executes multiple operations in parallel. As shown in the following example (Figure 24), multiple backup processes for all compute nodes occur in parallel.

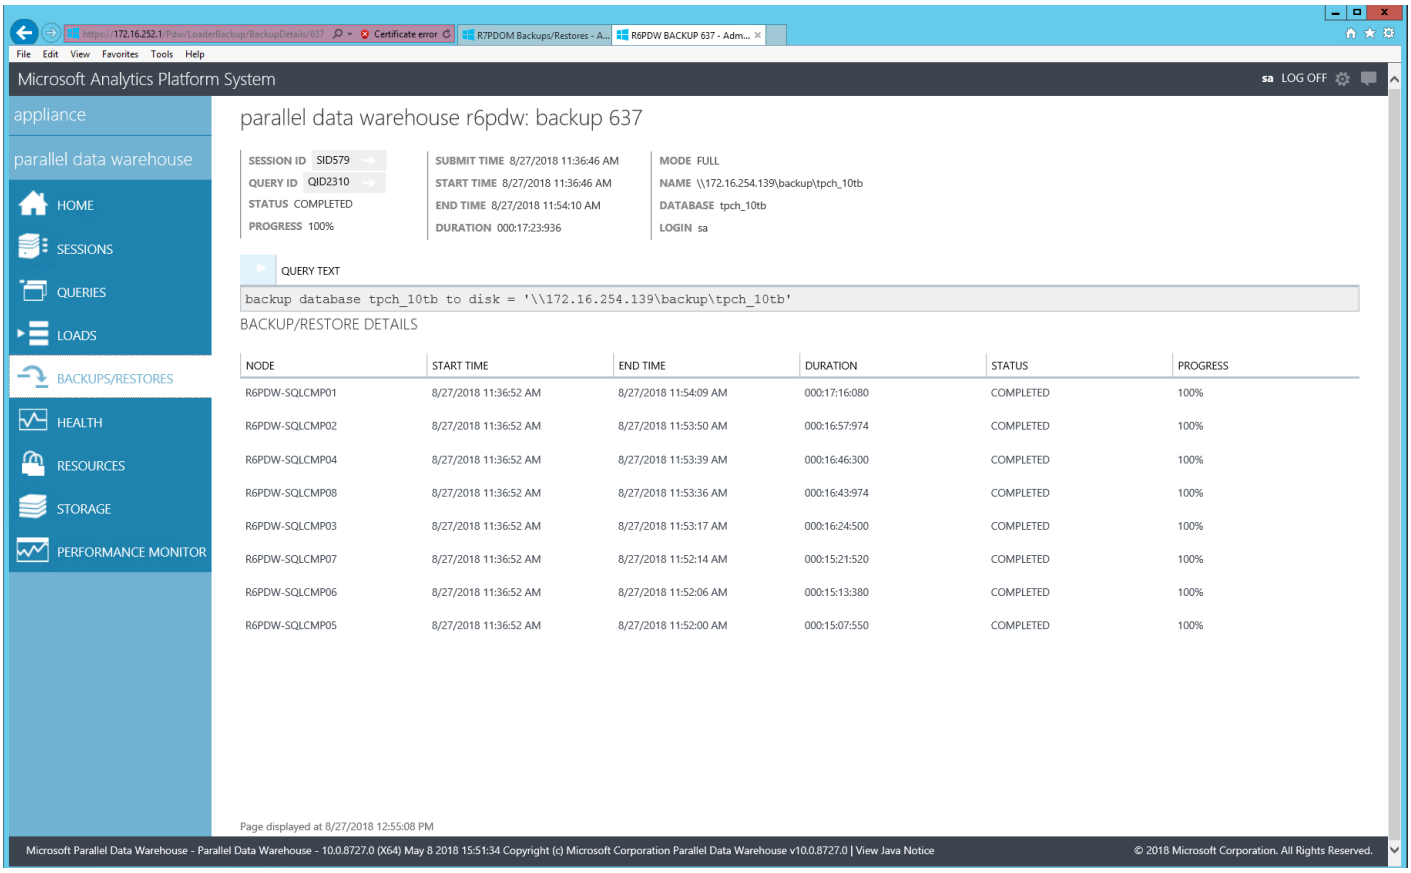

FIGURE 24. Parallel Data Warehouse (PDW) Backup Details

Health

As the default landing page, the Health Status screen (Figure 25) provides access to the overall status of PDW compute and control virtual machines. Clicking Alerts, All Alerts, or Errors will provide alerts or errors for all virtual machines. Selecting an individual node name will display additional details about the health status of that virtual machine.

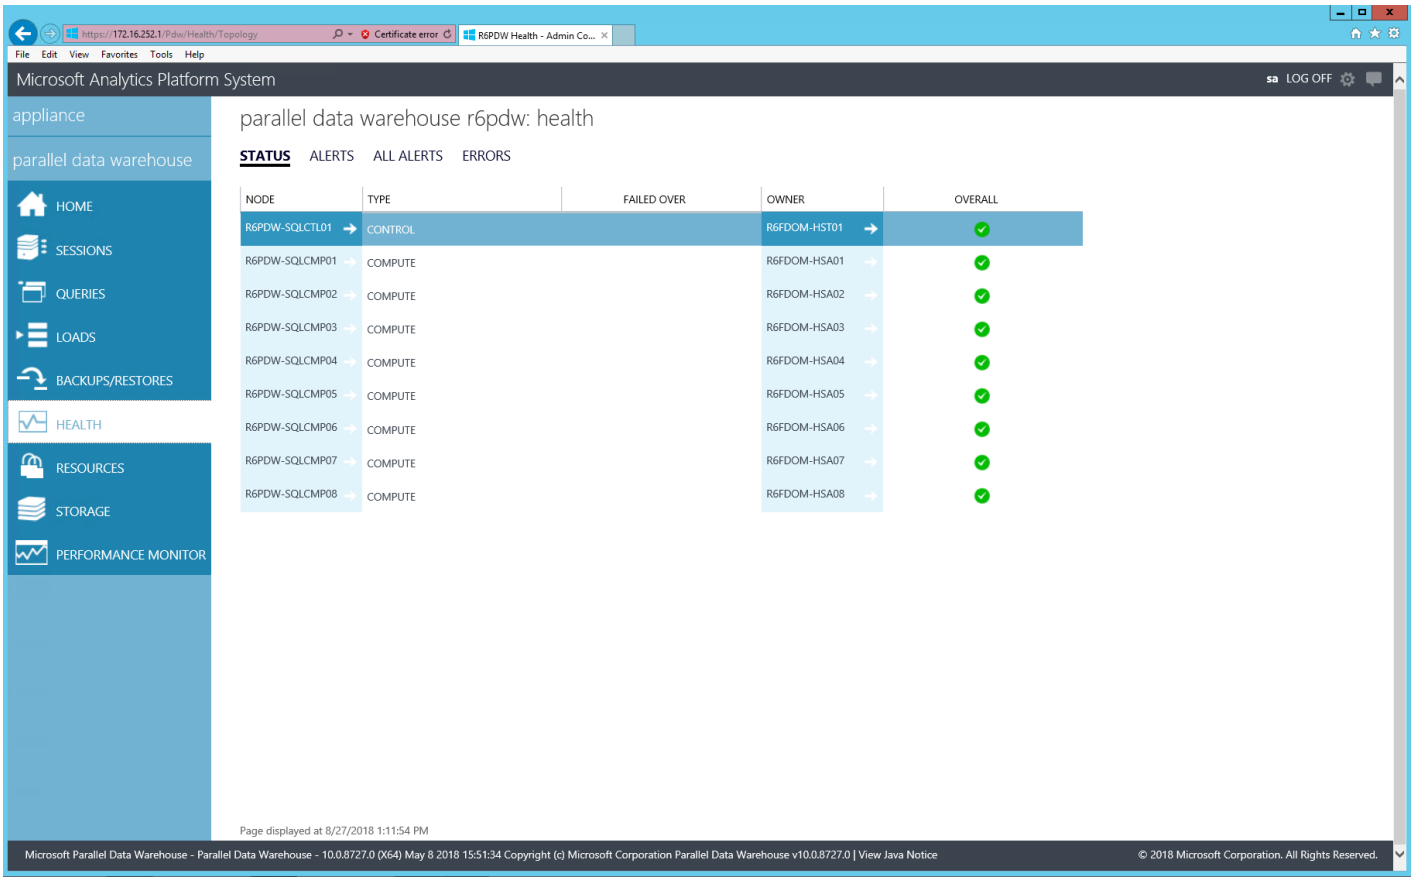

FIGURE 25. Parallel Data Warehouse (PDW) Health

Resources

The Resources screen (Figure 26) provides information about PDW resource Locks and Waits and their current status, which can affect the PDW resources and workload performance.

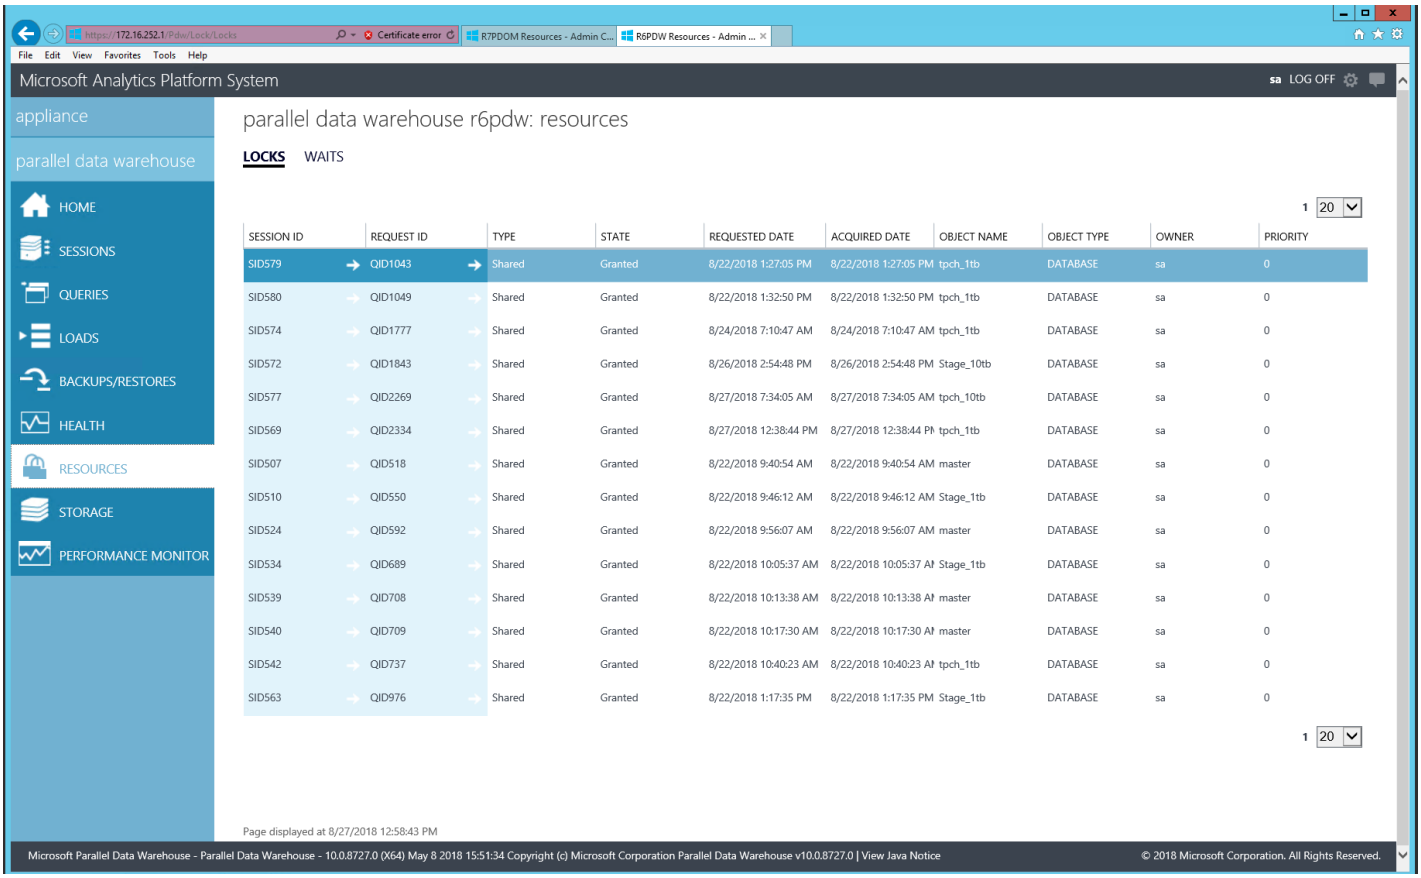

FIGURE 26. Parallel Data Warehouse (PDW) Resources

Storage

The Storage screen (Figure 27) provides a report showing the allocation of disk storage across various databases, including tempdb. This view provides information about the overall space utilization of PDW, including the Operating System and databases, and lists the databases by size (DATA (GB)).

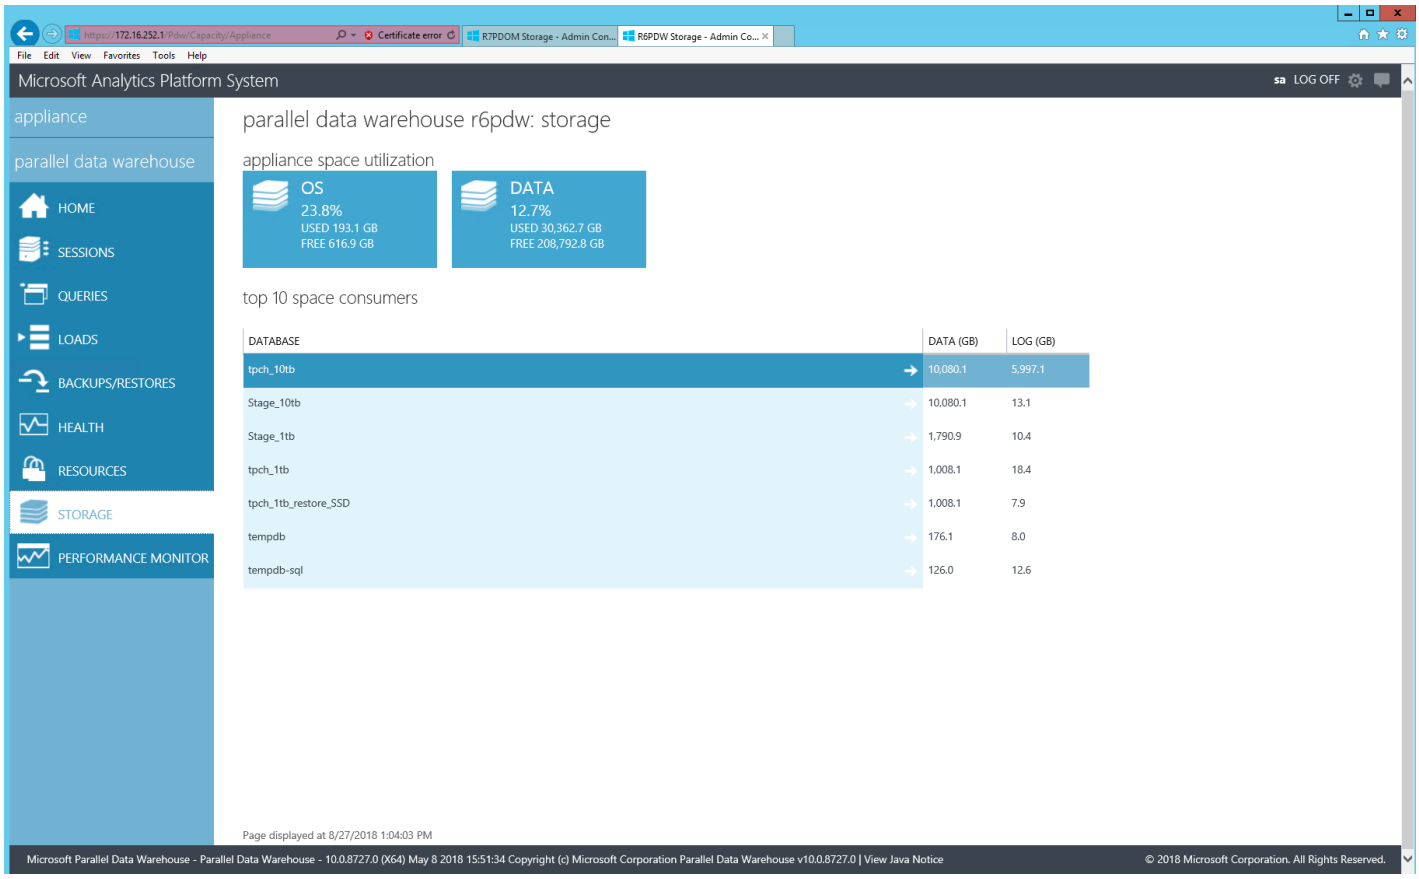

FIGURE 27. Parallel Data Warehouse (PDW) Storage Usage

Clicking a database name will show detailed information about that database. The following example (Figure 28) shows a 10TB database with data evenly distributed across all compute nodes.

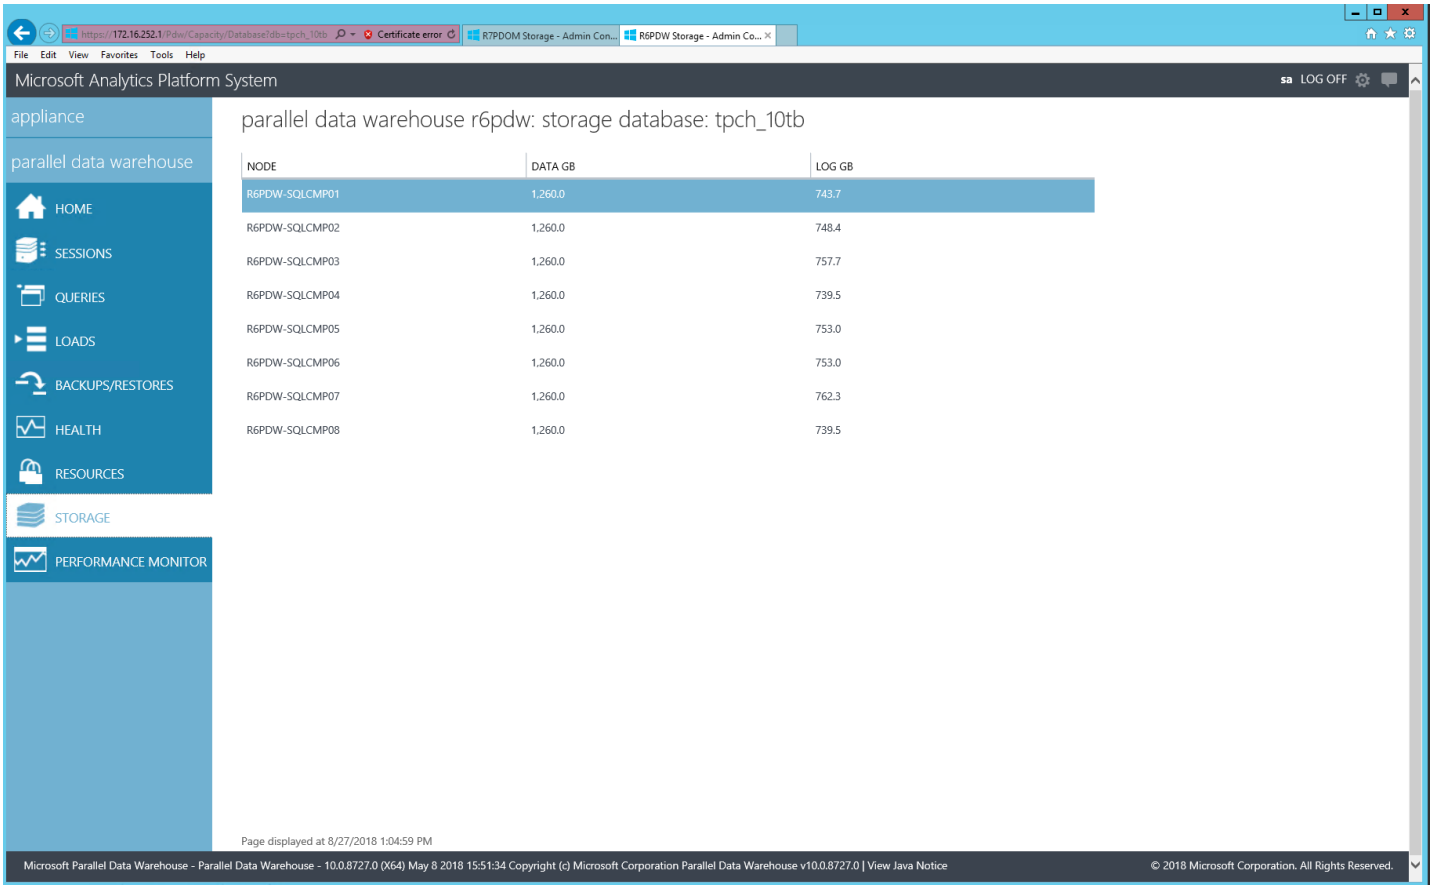

FIGURE 28. Parallel Data Warehouse (PDW) Storage Distribution

Performance Monitor

Use the Performance Monitor screen (Figure 29) to view the current performance in real time. Various selectable graphs (Figure 30) are available for users to monitor various performance counters for logical disks, memory, network interface, process, and processor.

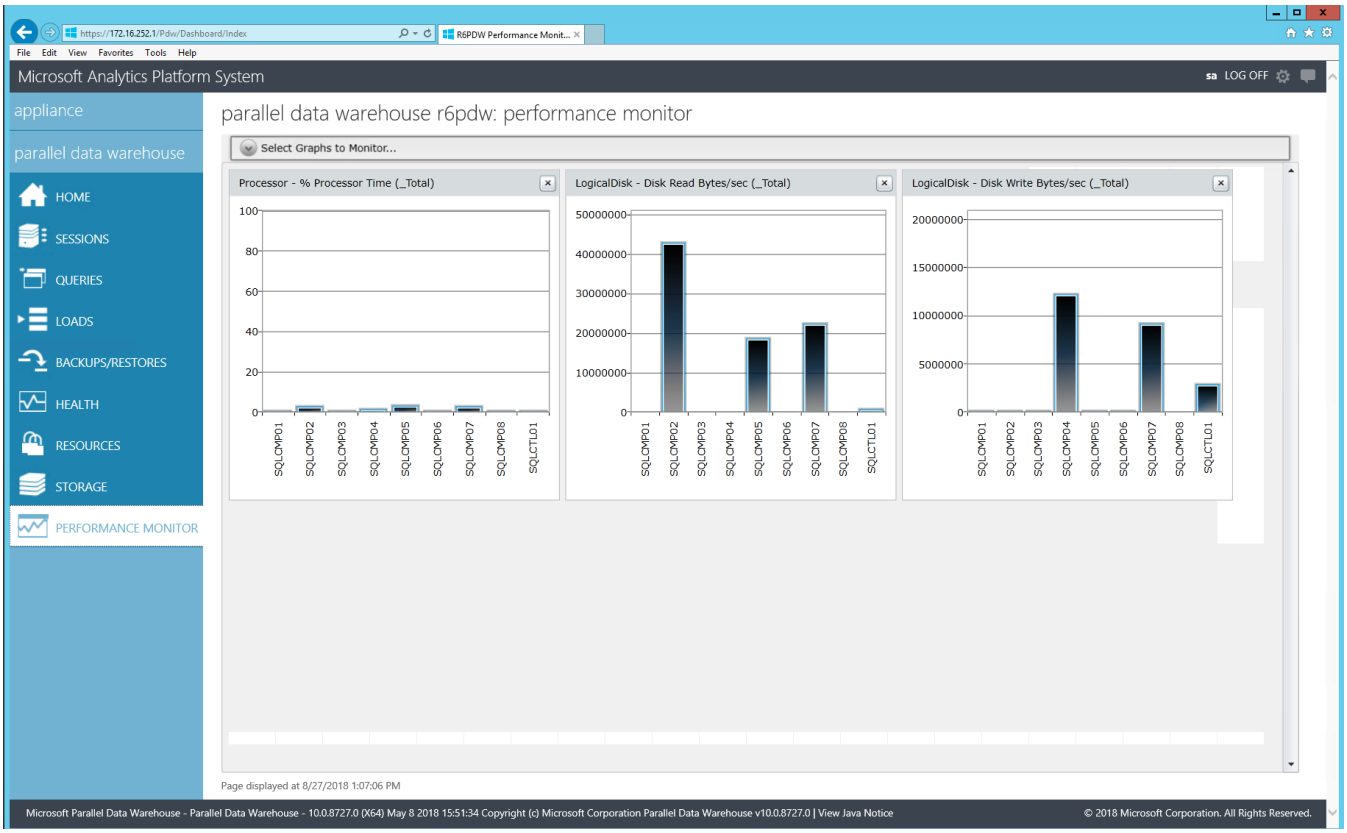

FIGURE 29. Parallel Data Warehouse (PDW) Performance Monitor

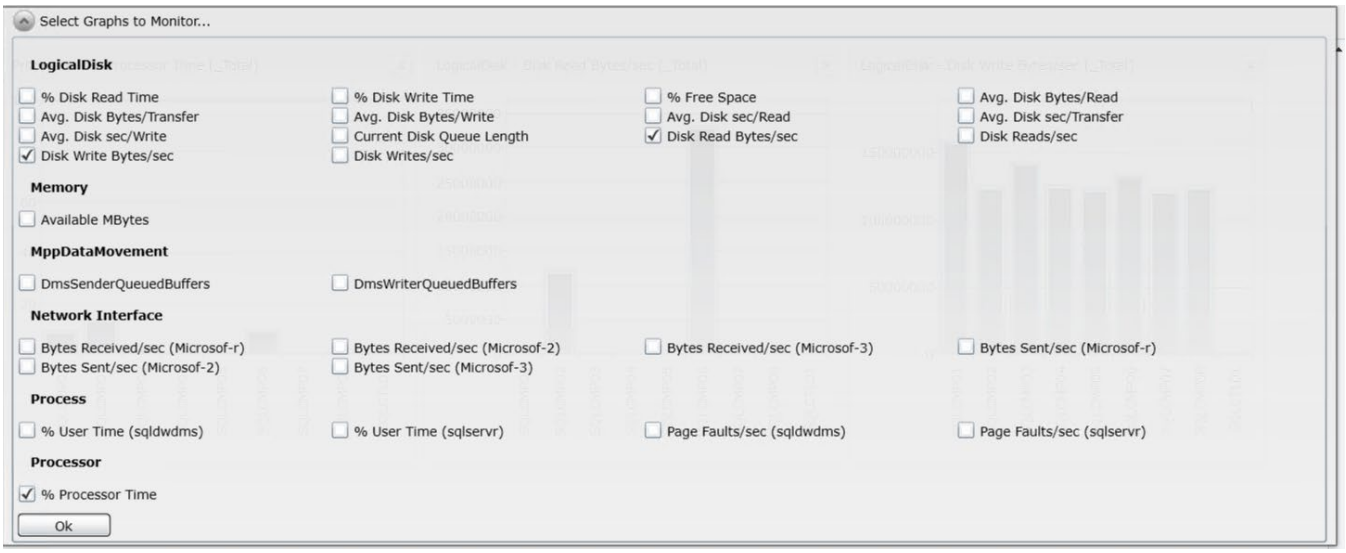

FIGURE 30. Parallel Data Warehouse (PDW) Selectable Graphs of Performance Monitor Counters

Windows Server management

The HPE ConvergedSystem 300 for Microsoft Analytics Platform utilizes clustered Windows virtual machines running on Microsoft Hyper-V® services. Management of the Windows virtual machines and clusters is performed through the use of the "Hyper-V Manager" and "Failover Cluster Manager" tools that are built into the Windows Server operating system.

Some common Hyper-V management tasks are included in Appendix B of this document. These are typical Hyper-V infrastructure tasks including:

- Add Hyper-V nodes to the infrastructure
- Cluster Roles alignment procedures – Tasks for ensuring that Hyper-V VMs are aligned to the correct Hyper-V server within the infrastructure.

Configuration Manager

Overview

The configuration of the solution is controlled through the Microsoft Analytics Platform System Configuration Manager, known as dwconfig.exe. This program resides on the CTL01 control node. This interface (Figure 31) is used for significant configuration changes, such as:

- Viewing the topology
- Changing administrative passwords for the components
- Installing certificates for validation of SSL connections to the web interfaces
- Configuring the Windows firewall rules
- Stopping and starting regions
- Restoring the master database
- Enable/disable features and settings

This tool will only be used by HPE ConvergedSystem 300 for Microsoft Analytics Platform administrators. This utility is also used by Hewlett Packard Enterprise during the initial installation process for basic configuration steps.

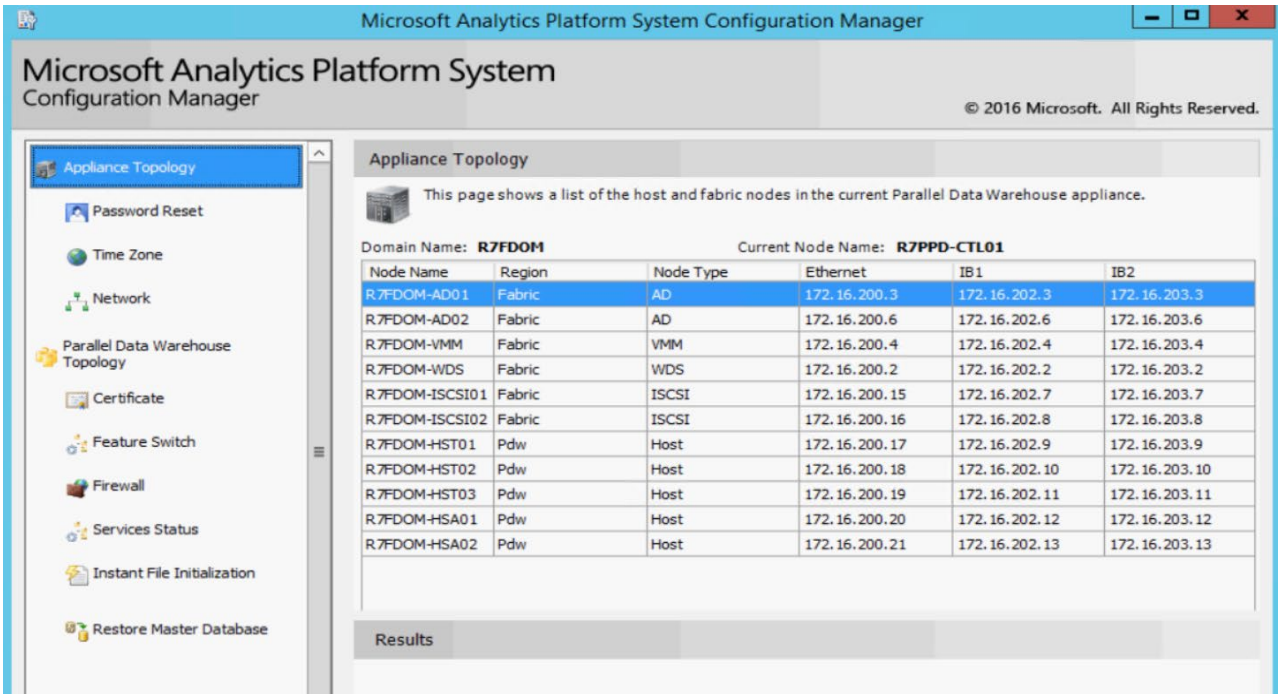

FIGURE 31. The Configuration Manager for the HPE ConvergedSystem 300 for Microsoft Analytics Platform

Running Configuration Manager

After logging onto CTL01 as the domain administrator, you can manage many PDW functions, as well as global configuration changes, using dwconfig.exe, the Configuration Manager. A shortcut to the utility will already be on the administrator's desktop.

From this application, you can view and modify many of the HPE ConvergedSystem 300 for Microsoft Analytics Platform configuration settings. After launching Configuration Manager, the "Appliance Topology" section is displayed (see figure 32). Several other configuration sections are available. For more information see, and Appliance Configuration.

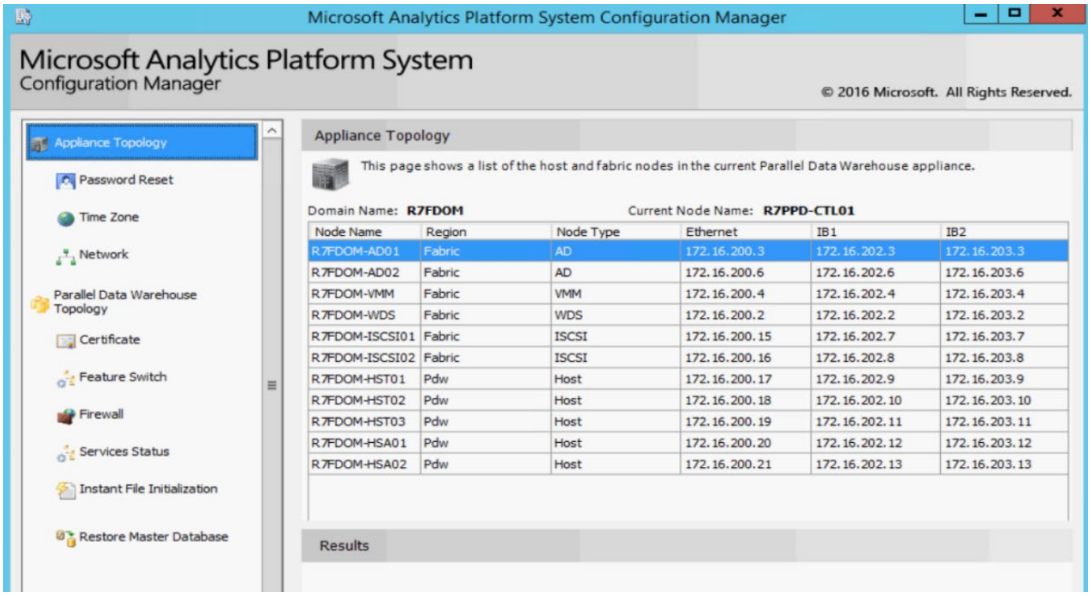

FIGURE 32. Configuration Manager Appliance Topology

Appliance Feature Switch

Use the Feature Switch page, as shown in figure 33, to update or enable/disable features and settings in Analytics Platform System. Changing feature switch values requires a service restart.

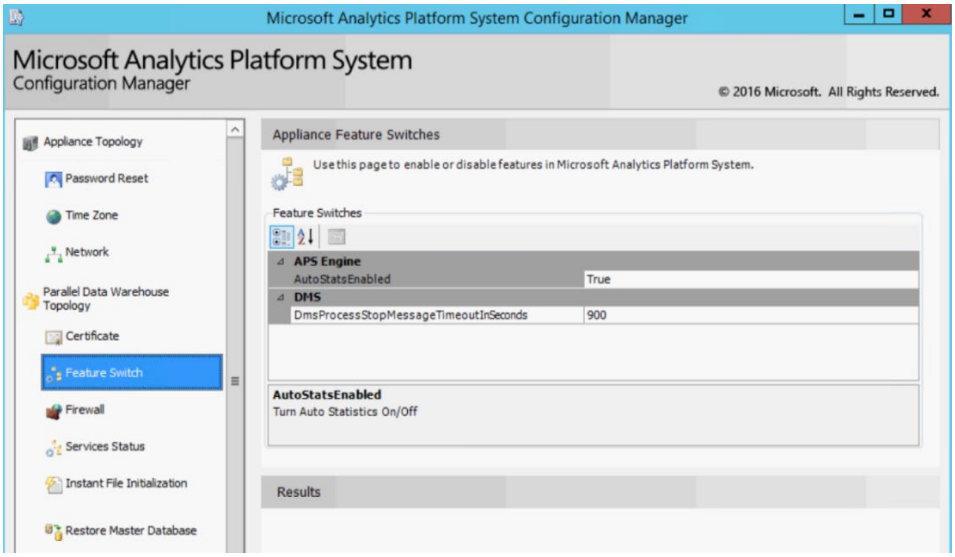

FIGURE 33. Configuration Manager. Appliance Feature Switches

The AutoStatsEnabled switch controls the auto statistics feature. This feature switch is set to true by default. Any database created after the upgrade will inherit auto creation and asynchronous update of statistics. For existing databases, database administrators can enable auto statistics with *alter database* command. For more information, see Statistics.

The DMSProcessStopMessageTimeoutInSeconds switch controls the time Data Movement Service (DMS) waits to synchronize on a busy system when a query involving data movement is cancelled. This value to 900 seconds (15 minutes) by default. The valid range is 0 – 3600 seconds.

DATA MANAGEMENT

Additionally, management of the Microsoft SQL Server Parallel Data Warehouse environment can be accomplished through a query interface. Command line interfaces, such as sqlcmd.exe are available. The primary graphical query environment is accomplished through the Microsoft SQL Server Data Tools (SSDT), an add-on to Microsoft Visual Studio 2012, and integrated into Microsoft Visual Studio 2013. Figure 34 shows the SQL Server Data Tools connected to a HPE ConvergedSystem 300 for Microsoft Analytics Platform, PDW region with Microsoft Visual Studio 2012.

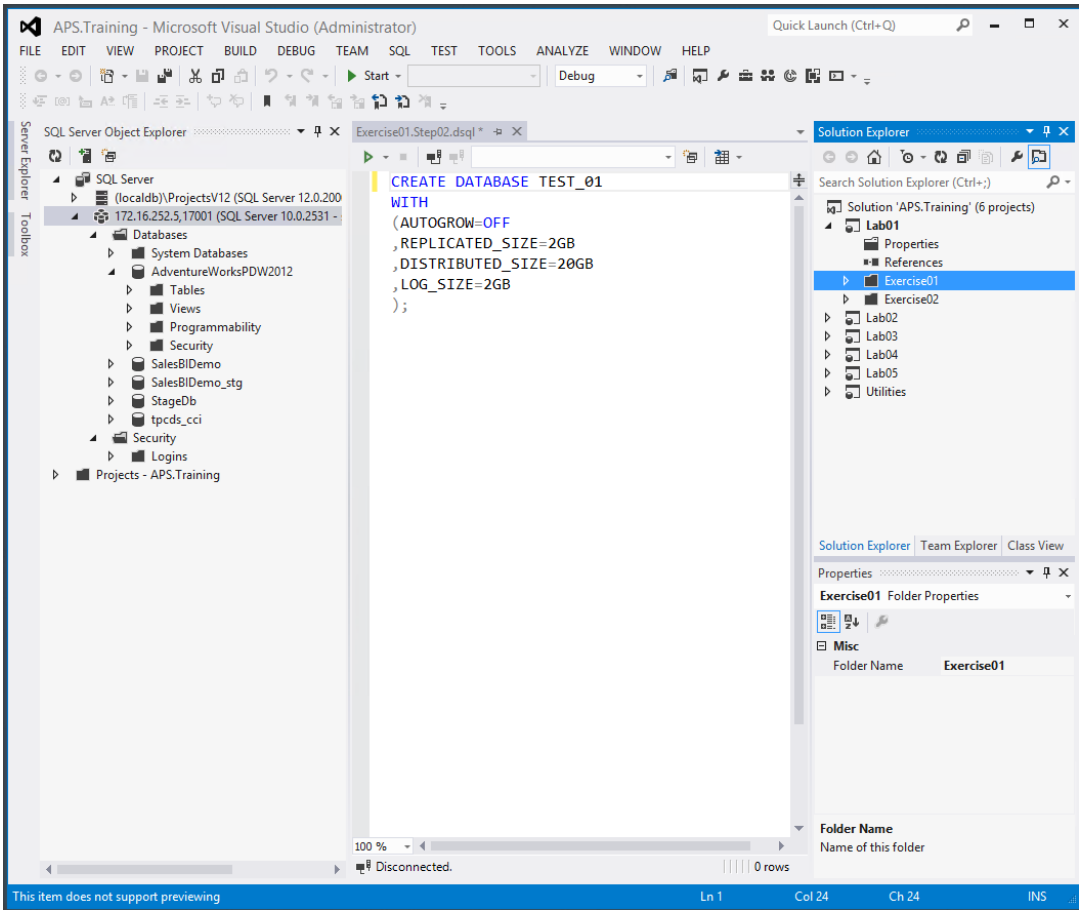

FIGURE 34. The Microsoft SQL Server Data Tools interface, using Microsoft Visual Studio 2012

Data loading prerequisites

With the hardware and software up and running, you can load data into the system. Before loading data, see "Guidance for designing Distributed and Replicated Tables")" This documentation provides information about SQL Server Parallel Data Warehouse terminology regarding distributed tables and replicated tables, and provides best practices for designing a data warehouse for the system.

In addition, Hewlett Packard Enterprise recommends that you have a good logical and physical database design. The following sections provide insight into design considerations. Contact Microsoft Services for additional information regarding optimal SQL Server Parallel Data Warehouse database design techniques.

Best practices for performance

For the best performance from SQL Server Parallel Data Warehouse, try eliminating all indexes except for Clustered Columnstore Indexes. Using indexes encourages disk head movement, and excessive disk head movement slows data streaming because of the excessive seek time required.

Indexing strategies, database design, or data loading optimizations using SQL Server Parallel Data Warehouse best practices is beyond the scope of this document; however, Microsoft Services can assist with these capabilities.

Distributed and replicated tables

A critical issue with SQL Server Parallel Data Warehouse database design involves defining replicated tables versus distributed tables. Tables default to replicated tables unless you define a distribution key. In general, large tables (such as fact tables) should be distributed, and small tables (such as dimension tables) should be replicated.

Figure 35 and Figure 36 show graphical representations of distributed and replicated tables.

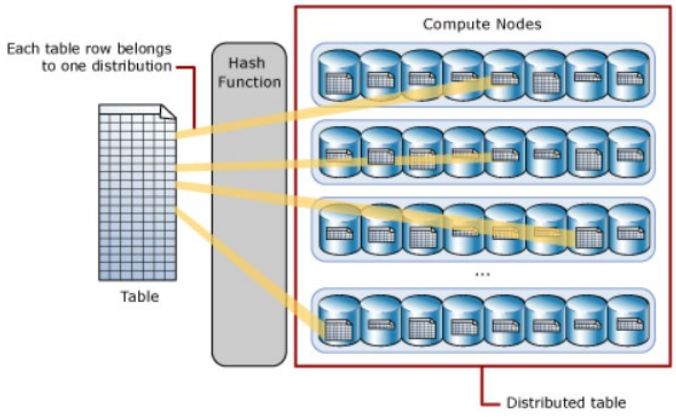

FIGURE 35. Distributed Tables

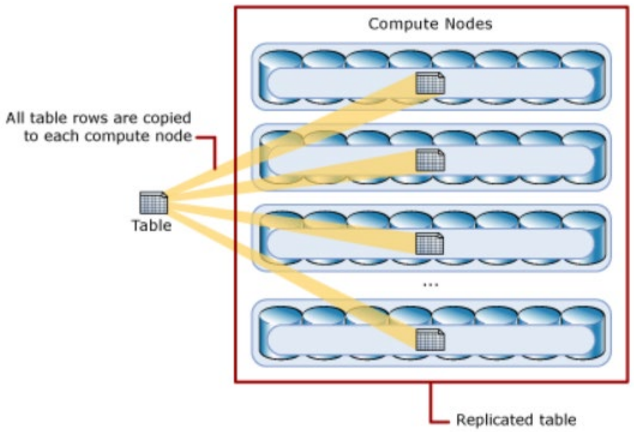

FIGURE 36. Replicated Tables

HPE ConvergedSystem 300 for Microsoft Analytics Platform running SQL Server Parallel Data Warehouse enables parallel execution of queries. The combination of hardware and software provides high levels of I/O throughput while maintaining consistent and predictable response times, which is useful when executing ad-hoc queries.

In-Memory analytics engine

The Microsoft SQL Server in-memory analytics engine, and the "memory-optimized" columnstore index are both features of the Microsoft Analytics Platform. System. The columnstore index feature, enables high performance query processing in SQL Server data warehouses and data marts. For example, business intelligence end users tend to submit queries that select a relatively small number of columns (hot columns) in the fact tables or dimension tables. This implies that, rather than having a database retrieve all the columns in a row, it makes sense to structure the database in a columnar fashion. Columnstore table structures only need to access, and potentially cache, the small percentage of columns that a majority of user's request.

Efficient compression is a side effect of a columnstore design in data warehousing. Database compression rates for data warehouses and data marts tend to be significantly greater when data is stored by column versus storing data by row. This is especially noticeable when a column has low cardinality.

By reorienting data storage in a column rather than a row, organizations can see HPE ConvergedSystem 300 for Microsoft Analytics Platform performance improve by reducing the query time from hours to seconds and compression ratios for data. In addition, Microsoft Parallel Data Warehouse enables updates to the HPE ConvergedSystem 300 for Microsoft Analytics Platform columnstore table.

Data loading tools

HPE ConvergedSystem 300 for Microsoft Analytics Platform does not require the purchase of a data loading server. Instead, Hewlett Packard Enterprise recommends a server with sufficient memory, storage, and connectivity (both Ethernet and dual-InfiniBand networking cards) be used as a platform for Extract, Transform and Load (ETL) and dwloader, HPE Insight Remote Support, and backup/restore operations. Because most site requirements do not promote backing up the HPE ConvergedSystem 300 for Microsoft Analytics Platform in tandem with a data upload, combining ETL and landing zone functionality on the same physical hardware as the backup makes sense.

Be sure to consider network connectivity when connecting a server for landing zone functions. For maximum throughput, Hewlett Packard Enterprise recommends InfiniBand, followed by 10Gb Ethernet. The slowest connection is 1Gb Ethernet connectivity between the data loading/backup node and the HPE ConvergedSystem 300 for Microsoft Analytics Platform.

There are two ways to load data into the SQL Server Parallel Data Warehouse database:

- Load data directly from SQL Server Integration Services (SSIS).

  If SSIS ETL software executes on an external server with the suggested InfiniBand connections, data can load directly from the server hosting SSIS into the staging database and into SQL Server Parallel Data Warehouse tables.

- Load data using dwloader (command line tool) from a cleansed and pre-formatted flat file.

  If you are loading data from a cleansed flat file residing on the suggested server, dwloader can perform faster than SSIS.

Hewlett Packard Enterprise has developed a reference guide for the Data Integration Platform (DIP). Tested configurations with complete Bills of Materials (BOMs) are provided. For more information, see Data Integration Platform - HPE ProLiant DL380 Gen10 reference guide

Loading data

SQL Server Parallel Database Warehouse data will typically be loaded into the appropriate database using one of the following mechanisms:

- ETL Software solution, such as SQL Server Integration Services (SSIS)
- dwloader utility (provided with the HPE ConvergedSystem 300 for Microsoft Analytics Platform solution)
- A custom data loading process incorporating one or more of the above tools
- A third party ETL solution

For most usages of data loading, an interim platform must be used to stage the data before it can be loaded into the SQL Server PDW database. Assuming, as above, that the Data Integration Platform is connected to the solution via the InfiniBand network, data can be transferred and loaded using the high bandwidth InfiniBand interconnects.

The Data Integration Platform is also an ideal location for any workflow or transformation engines needed to support the data loading activities. Note that most data loading activities will not necessarily require significant CPU scalability, but the ETL software may be more CPU intensive.

If being used as a data staging location, focus on the storage capabilities and capacity of the Data Integration Platform storage system. If you are also planning on hosting your ETL engine (such as SSIS) on the server, consider adding additional CPU and memory capacity to the server.

Backup/Restore

A key requirement of most database and analytics data storage systems is to have a backup of the data stored on the server. This backup location is typically going to be migrated off-site for extended storage and disaster recovery purposes. However, in the interim, a local copy of the backup is taken for efficiency and performance of the backup itself. The Data Integration Platform is the ideal location as the first storage location for HPE ConvergedSystem 300 for Microsoft Analytics Platform backups.

SQL Server Parallel Data Warehouse backup

SQL Server Parallel Data Warehouse (PDW) offers two backup types for your PDW databases. You can either take a full database backup (as it sounds, a copy of everything), or a differential backup, which will contain everything in your database that has changed since the last time a full database backup was taken. A differential backup will depend upon the last full backup for its baseline. For more information, review the Backup and Restore.

For each of these backup types, SQL Server Parallel Data Warehouse produces a set of files. These files together, rather than a single file, make up the database backup or differential backup. These backups will always use compression. The BACKUP DATABASE syntax assumes that the backup location will be on disk, via a UNC path to a specific directory to retain the backup (and must NOT be on the system itself). Therefore, a separate backup server is needed, with enough disk storage space to contain the full database backups and differential backups for each user database. Additionally, as with any local backup strategy an off-site or enterprise backup solution should be employed to make secondary (off-site) backups. Also keep in mind that a backup cannot be restored to a smaller system, so if a disaster recovery strategy involves using the backups, the other system must have at least the same number of compute nodes.

You should also back up the master database of the control node. Note, however, that a backup of the master database is a regular SQL Server backup of the master database, and must always be a full database backup. The master database will likely not be significant in size, but is an integral part of any recovery operation. It is strongly recommended that the master database backups use a parallel strategy to user database backups in terms of off-site storage, DR recovery planning, etc.

SQL Server Parallel Data Warehouse restore

To restore any individual database, a copy of the last full database backup must be available at an accessible UNC location. Once the full database backup is applied, the latest differential backup can then be applied to bring the database up to the latest point of recoverability. The last differential backup must also be accessible via a UNC share for the restore process.

If the entire system must be restored, the restore of the master database on the control node should be performed, and then the recovery as noted above for each individual user database (via the full and latest differential backup). Each of these backup file sets should be present on the Data Integration Platform for optimal recovery performance.

Using the Data Integration Platform

The recommendation for the backup and restore platform is that the server should be connected to the solution on the InfiniBand network if possible, to maximize the performance of the data movement, and use 10Gb Ethernet to connect to the customer corporate network. The Data Integration Platform as documented in the HPE ConvergedSystem 300 for Microsoft Analytics Platform: Data Integration Platform - HPE ProLiant DL380 Gen10 reference guide, meets these requirements and will provide optimal networking and data transfer capabilities for the solution.

Data archival

It is strongly recommended that the database backups will be copied to a remote storage location (for example, a Disaster Recovery data center). The Data Integration Platform is an ideal location to have the “latest” backups (for faster recovery purposes), but typically does not have enough storage to have multiple copies of database backups. Additionally, for disaster recovery planning, a remote copy of the data is critical in the event of a severe event affecting the primary data center.

HPE CONVERGEDSYSTEM 300 FOR MICROSOFT ANALYTICS PLATFORM LICENSING

When you purchase a HPE ConvergedSystem 300 for Microsoft Analytics Platform it is delivered as a pre-built appliance with software, hardware, and networking components already pre-installed and configured.

Use of Microsoft Windows Server 2012 R2, Microsoft SQL Server Parallel Data Warehouse Edition, and the Virtual Machine Manager component of Microsoft System Center 2012 Standard SP1 on this server hardware, is subject to the terms and conditions of the Volume License program agreement under which you have acquired licenses to these Microsoft products. You may not use the above software without first acquiring valid server and/or client access licenses from Microsoft or its authorized distributor or reseller. SQL Server Enterprise Per-Core licenses with Software Assurance may be used to license and run SQL Server PDW on the Analytics Platform System appliance. For more information, see Analytics Platform System Licensing.

SUPPORT RESOURCES

HPE ConvergedSystem 300 for Microsoft Analytics Platform Support Pack

Keeping your HPE ConvergedSystem 300 for Microsoft Analytics Platform running at maximum performance is critical and requires updating firmware and drivers to maintain the overall health of servers, storage arrays, switches, and all components included in the hardware stack. Developing a methodology for maintaining firmware on a complex hardware stack like the multi-server, MPP system powered by SQL Server Parallel Data Warehouse Edition, and HPE ConvergedSystem 300 for Microsoft Analytics Platform, requires a number of considerations.

HPE ConvergedSystem 300 for Microsoft Analytics Platform requires a specific, tested firmware and driver software stack that uses consistent firmware and driver versions across all similar nodes and subsystems throughout the system. Designed to help you simplify installation and maintenance tasks, the HPE ConvergedSystem 300 for Microsoft Analytics Platform Support Pack utility suite provides a unique set of tools and a collection of solution-tested HPE firmware and drivers to maintain and support the correct Microsoft Reference Architecture (MRA) configuration level. All maintenance procedures require careful coordination with support.

Hewlett Packard Enterprise support customers can download the HPE ConvergedSystem 300 for Microsoft Analytics Platform Support Pack from the HPE Software Updates website (hpe.com/downloads/software).

The support pack provides useful maintenance utilities, assists with support events, and generates needed reports for the Proactive Care Team.

Key components of the HPE ConvergedSystem 300 for Microsoft Analytics Platform Support Pack include:

- Tested and validated Firmware/Driver package
- A validation utility to confirm the system is in a supported configuration
- Diagnostics tools for deeply integrated health checks
- Reporter utility for deep device discovery and reporting
- Proactive Care Report Utility for a dashboard of the system configuration

These utilities are all value-add capabilities only available with the HPE ConvergedSystem 300 for Analysis Platform solution.

Download and installation

1. Download the latest Support Pack from the Hewlett Packard Enterprise website (hpe.com/tsusbportal/index.do?lc=EN_US).
2. Log in using HPE Passport.
3. Enter the appropriate customer information under “Directly enter an SAID”.
4. Review and accept the Terms and Conditions.
5. Click View available products. If you have used this site before, click View previously selected products, and then skip the following step.
6. Enter APS in the Search for products in all categories / product centers field.
7. Click “Get software updates”.
8. Select the latest version of the Complete Tool, which includes all of the tools in the latest toolkit or select an individual Support Pack tool from the same location by choosing a tool from the list.
9. Copy the Support Pack zip file to the HST02 server.

10. Double-click the zip file, and then click Extract all files.
11. Unzip the appropriate hardware-based folder. Make note of the directory location of the files associated with the support pack so that you can run the various scripts provided with the tool.

HPEApplianceXmlGenerator component

- Creates the HPEAppliancedetails.xml file based off the current hardware configuration. This file is needed for the support pack to function.
- Maps software IPs from definition files
- Discovers Ethernet switch IPs, IB switch IPs and iPDU IPs
- Determines server models
- Determines IB switch type
- Determines topology and identifies existing PDW regions
- Detailed logs are created for easy troubleshooting
- Supports v1 and v2 hardware

Built-In features

Validator features

- By default, calls HPEApplianceXmlGenerator
- Verifies Windows PowerShell is run with elevated privileges and user is in domain admin group
- Checks topology
- Checks server models
- Checks IB switch type
- Checks switch IPs (Ethernet/InfiniBand)
- Checks all bare metal hosts and VMs
- Pings components
- Verifies and validates common password and domain password
- Validates firmware/driver compliance against MRA for all components
- Iterates through all MRA files and finds closest MRA matching recipe
- Reports firmware/driver version mismatches
- Creates detailed logs of all operations and errors

NOTE

Customers are advised to upgrade the firmware and driver as reported by Validator tool. The 'HPE CS 300 APS Upgrade User Guide v2.7.7 650950-122.pdf' covers details for customers to upgrade the appliance as per current MRA. Disk drive firmware updates are not covered by MRA and customers are highly recommended to keep disk drive firmware up to date as per latest release of SPP. Refer 'Appendix E: Steps for updating SAS Drive firmware using custom SPP' from 'HPE CS 300 APS Upgrade User Guide v2.7.7 650950-122.pdf'

Diagnostics features

- By default, calls HPEApplianceXmlGenerator
- Verifies PowerShell is run with elevated privileges and user is in domain admin group
- Checks topology
- Checks server models

- Checks IB switch type
- Checks switch IPs (Ethernet/InfiniBand)
- Checks all bare metal hosts and VMs
- Pings components
- Verifies and validates common password and domain password
- Checks cable wiring (FC, iPDU, Ethernet and InfiniBand)
- Verifies storage including internal as well as external varieties of P2000/D6020. Hard disk drive capacities, health statuses, RAID levels, and logical volumes are all verified
- Verifies processors installed (processor models and speeds)
- Verifies memory installed (DIMM slots, sizes and amounts)
- Verifies PCIe components installed, along with which slots (FDR and HBA cards) into which they have been installed
- Verifies server configuration settings towards BIOS configuration settings
- Verifies Fibre channel/Ethernet/InfiniBand switch health, including fan, power, temperature status
- Creates detailed logs of all operations and errors

Reporter features

- By default, calls HPEApplianceXmlGenerator
- Verifies PowerShell is run with elevated privileges and user is in domain admin group
- Checks topology
- Checks server models
- Checks IB switch type
- Checks switch IPs (Ethernet/InfiniBand)
- Checks all bare metal hosts and VMs
- Pings components
- Verifies and validates common password and domain password
- Creates reports in csv and html file formats
- Report type 1: Serial number discovery for switches, storage enclosures, servers and iPDUs
- Report type 2: Firmware and driver version numbers for all components
- Report type 3: Physical hard disk drive properties (server, location, firmware/driver version, serial number, model)
- Creates detailed logs of all operations and errors

Proactive Care report features

- By default, calls HPEApplianceXmlGenerator
- Verifies PowerShell is run with elevated privileges and user is in domain admin group
- Checks topology
- Checks server models
- Checks IB switch type
- Checks switch IPs (Ethernet/InfiniBand)

- Checks all bare metal hosts and VMs
- Pings components
- Verifies and validates common password and domain password
- Provides a comprehensive rollup from the results of the Validator, Diagnostics and Reporter tools into a csv file (viewable from Excel)
- Report is sorted by component (rack ID, U location, device, IP address, serial number, make and model, status)
- Compiles serial numbers for hardware components (switches, iPDUs, servers, and storage enclosures)
- Report contains each component's validation issues with firmware and driver versions
- Report contains each component's diagnostic issues
- Creates detailed logs of all operations and errors

HPE Insight Remote Support software and HPE ConvergedSystem 300 for Microsoft Analytics Platform

HPE Insight Remote Support (Insight RS) is a key component of HPE Proactive Care support from Hewlett Packard Enterprise. This software provides continuous proactive monitoring of the hardware components within the HPE ConvergedSystem 300 for Microsoft Analytics Platform environment.

Your Hewlett Packard Enterprise representative installing the HPE ConvergedSystem 300 for Microsoft Analytics Platform will work with you to install and configure an Insight RS system. This system cannot be part of the HPE ConvergedSystem300 for Microsoft Analytics Platform.

For additional information, see about Insight Remote Support.

SERVICES

Included with the HPE ConvergedSystem 300 for Microsoft Analytics Platform, is factory integration, installation and support. This section outlines services that are available for HPE ConvergedSystem 300 for Microsoft Analytics Platform.

Services available for HPE ConvergedSysytem300 for Microsoft Analytics Platform

Support for the HPE ConvergedSystem300 for Microsoft Analytics Platform from the choices outlined below.

HPE Proactive Care – 3, 4 or 5 Year (3 year is included in purchase price)

HPE Proactive Care for the ConvergedSystem 300 for Microsoft Analytics Platform begins with providing all of the benefits of proactive monitoring and reporting along with access to Hewlett Packard Enterprise specially trained people to provide a complete solution-level support experience and put in place the fundamentals needed for stability and availability of the Microsoft Analytics Platform environment. Proactive Care helps in problem prevention, with predictive analytics, personalized analysis with recommendations and advice paired with rapid access to technical experts to help rapidly resolve any problem. You receive an enhanced call experience with advanced technical expertise, a single point of contact for the support of all components and end-to-end case ownership. You also benefit from the specially aligned, collaborative reactive processes between Hewlett Packard Enterprise and Microsoft.

Customers can customize their reactive support level by selecting either a 6-hour call-to-repair or 24x7 with 4-hour onsite response.

HPE Proactive Care Advanced – 3, 4 or 5 year

HPE Proactive Care Advanced is the recommended support for ConvergedSystem 300 for Microsoft Analytics Platform environments. It builds on HPE Proactive Care, providing additional benefits such as the assignment of a dedicated, local account support manager (ASM) for collaboration, best practices and critical event management that provides 24x7 fast response and IT service restoration with incident follow-up to prevent a repeat. All this this is designed to give you an incredibly personalized, high-touch support experience that keeps your system fully available and running at peak performance.

HPE Proactive Care Advanced includes credits that you can use to select and fund the specialized service assistance you need, when you need it. You can choose from a range of predefined technical services on the HPE Proactive Select menu, or your ASM can work with you to define the specific advice or assistance you need. It may be for a specific project or to simply purchase the 'HPE Firmware Update Service for ConvergedSystem 300 for Microsoft Analytics Platform'.

HPE Datacenter Care (recommended by HPE)
HPE Datacenter Care Service is HPE's most comprehensive support solution tailored to meet your specific data center support requirements. It offers a wide choice of proactive and reactive service levels to cover requirements ranging from the most basic to the most business-critical environments. A mutually agreed upon and executed Statement of Work (SOW) will detail the precise combination of reactive and proactive support features that will be provided under HPE Datacenter Care Service based upon your requirements.

The service includes an assigned account team led by a trained Hewlett Packard Enterprise Account Support Manager (ASM). The team's goal is to form a close working relationship with designated members of your IT staff and gain a clear understanding of your business objectives, key service-level agreements (SLAs), and the key performance indicators (KPIs) you need to meet. Delivery of the various support options you have chosen will be overseen by the ASM and directed at meeting your goals.

For more information on these support offerings, refer to your Account Manager or Reseller.

Factory Express

Factory Express services are utilized by the HPE ConvergedSystem 300 for Microsoft Analytics Platform solution. Benefits include:

- **Business Agility**
  - Broad portfolio
  - Faster technology transactions
  - Smoother solution deployment
- **Managing Costs**
  - Save time, money and resources
  - Collaborate with Hewlett Packard Enterprise experts
  - One-touch efficiency
- **Mitigating Risk**
  - Predictable hassle-free experience
  - Efficiency and quality
- **Increase performance and quality**
  - Innovation based on standards

TASKS AND PROCEDURES FOR GETTING STARTED WITH THE HPE CONVERGEDSYSTEM300 FOR MICROSOFT ANALYTICS PLATFORM

Several common tasks and procedures for getting started with the HPE ConvergedSystem 300 for Microsoft Analytics Platform are listed in Appendix B. These procedures include:

- Initial DB setup and load – For creating initial staging, production and test databases on the appliance as well as recommendations on using various load procedures.
- Backing up and restoring databases – For creating a Backup and Restore plan for the appliance
- Adding users to the appliance – Procedures for adding new logins to the appliance
- Modifying device passwords – Several devices contain passwords that can be modified to conform to conventions or customer specs. These include, Ethernet and InfiniBand switches, PDUs, server iLO logins, etc.

For more information on details for other tasks, see Appliance management tasks for Analytics Platform System.

APPENDIX A: CONFIGURATION DETAILS

Solution hardware

TABLE A1. Hardware and software configurations

Component	Quantity	Description
Orchestration and Failover Servers	Minimum 2	HPE ProLiant DL360 Gen10 server, each with: • 2 x Intel Xeon-Gold 6134 processor • 256GB RAM (8 HPE 32GB PC4-2666V (DDR4-2666)) memory • 2 x SAS hard drives, 600GB/10,000 rpm 2.5 inch • HPE InfiniBand FDR/Ethernet 10Gb/40Gb 2-port 544+QSFP adapter • HPE Ethernet 1Gb 4-port 331FLR Adapter
Data Scale Unit (Server)	In groups of two	HPE ProLiant DL360 Gen10 server, each with: • 2 x Intel Xeon-Gold 6134 processor • 256GB RAM (8 HPE 32GB PC4-2666V (DDR4-2666)) memory • 2 x SAS hard drives, 600GB/10,000 rpm 2.5 inch • Smart Array208e-p SR for D6020 connectivity • HPE InfiniBand FDR/Ethernet 10Gb/40Gb 2-port 544+QSFP adapter • HPE Ethernet 1Gb 4-port 331FLR adapter
Data Scale Unit (Storage) – HPE D6020 disk enclosure	One per scale unit	HPE D6020 disk enclosure, each with: • Dual I/O Module option kit • 70 x 2TB HPE SAS 7.2K 3.5 inch HDD or 70 x 6TB HPE SAS 7.2K 3.5 inch HDD
Network switches	2	HPE 5900A48G 4XG 2QSFP+ Switch
InfiniBand switches	2	Mellanox InfiniBand EDR v2 100 Gb/sec 36-port Power-side-inlet Airflow Managed Switch
Infrastructure	1 per rack	HPE 42U 600x1075 G2 Enterprise Shock Rack, each with: • Side Panel Stabilizer Kit • Rack Grounding Kit • Single-phase (4 PDUs per rack) or Three-phase (2 PDUs per rack) power solution • Cabling and wiring to connect all components

Power diagrams

Figures A2 and A3 show an example of the power cabling that happens within a standard HPE ConvergedSystem 300 for Microsoft Analytics Platform rack, showing the redundancy and reliability of the design.

Facing the rear of the rack, PDU 2 and PDU 4 are on the left, PDU 1 and PDU 3 are on the right.

Single Phase Extension Bar Cabling - REAR VIEW

Scroll down to see three phase cabling.

Rack 1:

PDU	Load Segment		Port	Device	RU-Loc
2	1	Extension Bar 1	1	R1IBSW1 Power Supply (Left)	U-42
			2	R1IBSW2 Power Supply (Left)	U-41
			3	R1ETHSW1 Power Supply (Left)	U-40
			4	R1ETHSW2 Power Supply (Left)	U-39
			5	HST01 Power Supply 2 (Left)	U-38
	2	Extension Bar 2	1	HST02 Power Supply 2 (Left)	U-37
			2	HST03 Power Supply 2 (Left)	U-36
			3	Not Used	
			4	Not Used	
			5	Not Used	
	3	Extension Bar 3	1	HSA08 Power Supply 2 (Left)	U-28
			2	HSA07 Power Supply 2 (Left)	U-27
			3	R1S4 Drawer2 Primary I/O (Top Left)	U-22
			4	R1S4 Drawer2 Secondary I/O (Btm Left	U-22
			5	Not Used	
	4	Not Used			
	5	Not Used			
	6	Not Used			

PDU	Load Segment		Port	Device	RU-Loc
1	1	Extension Bar 1	1	R1IBSW1 Power Supply (Right)	U-42
			2	R1IBSW2 Power Supply (Right)	U-41
			3	R1ETHSW1 Power Supply (Right)	U-40
			4	R1ETHSW2 Power Supply (Right)	U-39
			5	HST01 Power Supply 1 (Right)	U-38
	2	Extension Bar 2	1	HST02 Power Supply 2 (Right)	U-37
			2	HST03 Power Supply 2 (Right)	U-36
			3	Not Used	
			4	Not Used	
			5	Not Used	
	3	Extension Bar 3	1	HSA08 Power Supply 1 (Right)	U-28
			2	HSA07 Power Supply 1 (Right)	U-27
			3	R1S4 Drawer1 Primay I/O (Top Right)	U-22
			4	R1S4 Drawer1 Secondary I/O (Btm Right)	U-22
			5	Not Used	
	4	Not Used			
	5	Not Used			
	6	Not Used			

PDU	Load Segment		Port	Device	RU-Loc
4	1	Not Used			
	2	Not Used			
	3	Not Used			
	4	Extension Bar 4	1	HSA06 Power Supply 2 (Left)	U-21
			2	HSA05 Power Supply 2 (Left)	U-20
			3	R1S3 Drawer2 Primary I/O (Top Left)	U-15
			4	R1S3 Drawer2 Secondary I/O (Btm Left	U-15
			5	Not Used	
	5	Extension Bar 5	1	HSA04 Power Supply 2 (Left)	U-14
			2	HSA03 Power Supply 2 (Left)	U-13
			3	R1S2 Drawer2 Primary I/O (Top Left)	U-8
			4	R1S2 Drawer2 Secondary I/O (Btm Left	U-8
			5	Not Used	
	6	Extension Bar 6	1	HSA02 Power Supply 2 (Left)	U-7
			2	HSA01 Power Supply 2 (Left)	U-6
			3	R1S1 Drawer2 Primary I/O (Top Left)	U-1
			4	R1S1 Drawer2 Secondary I/O (Btm Left	U-1
			5	Not Used	

PDU	Load Segment		Port	Device	RU-Loc
3	1	Not Used			
	2	Not Used			
	3	Not Used			
	4	Extension Bar 4	1	HSA06 Power Supply 1 (Right)	U-21
			2	HSA05 Power Supply 1 (Right)	U-20
			3	R1S3 Drawer1 Primay I/O (Top Right)	U-15
			4	R1S3 Drawer1 Secondary I/O (Btm Right)	U-15
			5	Not Used	
	5	Extension Bar 5	1	HSA04 Power Supply 1 (Right)	U-14
			2	HSA03 Power Supply 1 (Right)	U-13
			3	R1S2 Drawer1 Primay I/O (Top Right)	U-8
			4	R1S2 Drawer1 Secondary I/O (Btm Right)	U-8
			5	Not Used	
	6	Extension Bar 6	1	HSA02 Power Supply 1 (Right)	U-7
			2	HSA01 Power Supply 1 (Right)	U-6
			3	R1S1 Drawer1 Primay I/O (Top Right)	U-1
			4	R1S1 Drawer1 Secondary I/O (Btm Right)	U-1
			5	Not Used	

FIGURE A2. Power cable wiring in an HPE ConvergedSystem 300 for Microsoft Analytics Platform – Single Phase power

Three Phase Extension Bar Cabling - REAR VIEW

Rack 1

PDU	Load Segment	Extension Bar 1		
		Port	Device	RU-Loc
2	1	1	R1IBSW1 Power Supply (Left)	U-42
		2	R1IBSW2 Power Supply (Left)	U-41
		3	R1ETHSW1 Power Supply (Left)	U-40
		4	R1ETHSW2 Power Supply (Left)	U-39
		5	HST01 Power Supply 2 (Left)	U-38
	2	Extension Bar 2		
		Port	Device	RU-Loc
		1	HST02 Power Supply 2 (Left)	U-37
		2	HST03 Power Supply 2 (Left)	U-36
		3	Not Used	
		4	Not Used	
		5	Not Used	
	3	Extension Bar 3		
		Port	Device	RU-Loc
		1	HSA08 Power Supply 2 (Left)	U-28
		2	HSA07 Power Supply 2 (Left)	U-27
		3	R1S4 Drawer2 Primary I/O (Top Left)	U-22
		4	R1S4 Drawer2 Secondary I/O (Btm Left	U-22
		5	Not Used	
	4	Extension Bar 4		
		Port	Device	RU-Loc
		1	HSA06 Power Supply 2 (Left)	U-21
		2	HSA05 Power Supply 2 (Left)	U-20
		3	R1S3 Drawer2 Primary I/O (Top Left)	U-15
		4	R1S3 Drawer2 Secondary I/O (Btm Left	U-15
		5	Not Used	
	5	Extension Bar 5		
		Port	Device	RU-Loc
		1	HSA04 Power Supply 2 (Left)	U-14
		2	HSA03 Power Supply 2 (Left)	U-13
		3	R1S2 Drawer2 Primary I/O (Top Left)	U-8
		4	R1S2 Drawer2 Secondary I/O (Btm Left	U-8
		5	Not Used	
	6	Extension Bar 6		
		Port	Device	RU-Loc
		1	HSA02 Power Supply 2 (Left)	U-7
		2	HSA01 Power Supply 2 (Left)	U-6
		3	R1S1 Drawer2 Primary I/O (Top Left)	U-1
		4	R1S1 Drawer2 Secondary I/O (Btm Left	U-1
		5	Not Used	

PDU	Load Segment	Extension Bar 1		
		Port	Device	RU-Loc
1	1	1	R1IBSW1 Power Supply (Right)	U-42
		2	R1IBSW2 Power Supply (Right)	U-41
		3	R1ETHSW1 Power Supply (Right)	U-40
		4	R1ETHSW2 Power Supply (Right)	U-39
		5	HST01 Power Supply 1 (Right)	U-38
	2	Extension Bar 2		
		Port	Device	RU-Loc
		1	HST02 Power Supply 1 (Right)	U-37
		2	HST03 Power Supply 2 (Right)	U-36
		3	Not Used	
		4	Not Used	
		5	Not Used	
	3	Extension Bar 3		
		Port	Device	RU-Loc
		1	HSA08 Power Supply 1 (Right)	U-28
		2	HSA07 Power Supply 1 (Right)	U-27
		3	R1S4 Drawer1 Primay I/O (Top Right)	U-22
		4	R1S4 Drawer1 Secondary I/O (Btm Right)	U-22
		5	Not Used	
	4	Extension Bar 4		
		Port	Device	RU-Loc
		1	HSA06 Power Supply 1 (Right)	U-21
		2	HSA05 Power Supply 1 (Right)	U-20
		3	R1S3 Drawer1 Primay I/O (Top Right)	U-15
		4	R1S3 Drawer1 Secondary I/O (Btm Right)	U-15
		5	Not Used	
	5	Extension Bar 5		
		Port	Device	RU-Loc
		1	HSA04 Power Supply 1 (Right)	U-14
		2	HSA03 Power Supply 1 (Right)	U-13
		3	R1S2 Drawer1 Primay I/O (Top Right)	U-8
		4	R1S2 Drawer1 Secondary I/O (Btm Right)	U-8
		5	Not Used	
	6	Extension Bar 6		
		Port	Device	RU-Loc
		1	HSA02 Power Supply 1 (Right)	U-7
		2	HSA01 Power Supply 1 (Right)	U-6
		3	R1S1 Drawer1 Primay I/O (Top Right)	U-1
		4	R1S1 Drawer1 Secondary I/O (Btm Right)	U-1
		5	Not Used	

FIGURE A3. Power cable wiring in an HPE ConvergedSystem 300 for Microsoft Analytics Platform – Three Phase power

APPENDIX B: HPE CONVERGEDSYSTEM 300 FOR MICROSOFT ANALYTICS PLATFORM PROCEDURES AND TASKS

Initial database setup and load

Creating a staging database

Microsoft recommends creating and using only one staging database for optimal SQL Server Parallel Data Warehouse performance.

NOTE

Do not create tables in the staging database. When you run the dwloader tool, it creates tables for you.

SQL Server Parallel Data Warehouse includes the tempdb database for internal use by the software to store temporary tables on local disks of the Control node. As a best practice, avoid using tempdb. Instead, create your own production databases for your environment.

The following example of a command file shows how you can create a database and the associated tables as the destination for the loaded data.

```
// Getting Started Guide - Sample DB and table creation
// Create target database for load test
CREATE DATABASE db100gb
WITH (
   AUTOGROW = ON,
   REPLICATED_SIZE = 50,
   DISTRIBUTED_SIZE = 100,
   LOG_SIZE = 5
)
;
// Create staging database
```

For more information, see Using a staging database in Parallel Data Warehouse (PDW).

SQL Server Parallel Data Warehouse uses a staging database to store data temporarily during the load process. By default, Parallel Data Warehouse uses the destination database as the staging database, which can cause table fragmentation.

To reduce table fragmentation, create a user-defined staging database. If rollback from a load failure is not a concern, you can use the fastappend loading mode to improve performance. The fastappend loading mode skips the temporary staging table and loads directly into the destination table.

```
CREATE DATABASE staging
WITH (
   AUTOGROW = ON,
   REPLICATED_SIZE = 40,
   DISTRIBUTED_SIZE = 90,
   LOG_SIZE = 4
);

USE db100gb;

// Create partsupp table
// partsupp table will be distributed using ps_partkey column

CREATE TABLE partsupp   ( ps_partkey      bigint,
ps_suppkey      bigint,      ps_availqty     integer,
ps_supplycost   decimal(15,2),
ps_comment      varchar(199))
 WITH (Distribution = hash(ps_partkey));
```

Loading data using dwloader

After creating the databases and tables, you can load your data.

SQL Server Parallel Data Warehouse supports either ETL or ELT data loading techniques. In either case, load your data into the HPE ConvergedSystem 300 for Microsoft Analytics Platform from a server that contains the input data file.

NOTE
SQL Server Parallel Data Warehouse allows dwloader to execute from multiple servers when those servers act as data loading sources. Data loading timelines improve when ETL software executes on multiple servers in parallel.

Initiate your data load using dwloader. You can find dwloader at the following location:

```
%SystemDrive%\Program Files\Microsoft SQL Server Parallel Data Warehouse\100\dwloader.exe
```

The following SQL command is an example of a load command using the common parameters listed below.

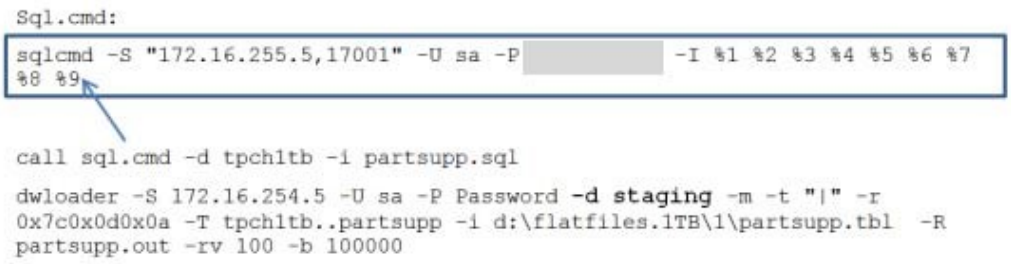

- `-S` – IP address of the ConvergedSystem300 for Microsoft Analytics Platform
- `-U` – User ID (in our example, sa).
- `-P` – System password.
- `-d` – Name of the staging database.

NOTE
Microsoft recommends using a staging database for optimal performance. A staging database helps Parallel Data Warehouse maintain sequential storage, which reduces disk seek times. The -d parameter specifies the use of the staging database. If you omit this parameter, data loads directly into the destination tables without using the staging database, which can reduce overall database performance.

- `-m` – Commit loads in parallel. Required option for `-M fastappend`. Performs much faster than the default loading mode by bypassing the staging database. Has no rollback function, so any recovery from a failed or aborted load must be handled by your own load process.
- `-t "|"` – Delimit each field (column) in the row.
- `-r 0x7c0x0d0x0a` – Delimit row.
- `-T` – Target database.
- `-i` – Source data location (input).
- `-R` – File where you can find the rows which failed to load.
- `-rv` – Number or percentage of row rejections to allow before halting the load (in this example, 100).
- `-b` – Batch size; the number of rows to load per transaction. Defaults to 10,000.

For more information about this command and its parameters, see dwloader Command-Line Loader for Parallel Data Warehouse.

After loading your data, you can use Client Tools that support OLE DB, ODBC, or ADO.NET to query your data. The most common tools used to query the database are Reporting Services, SQL Server Analysis Services (OLAP cubes), SQL Server Data Tools (SSDT), and SSIS.

SSIS is useful when you need to perform a table look-up during ETL processing. SQL Server Analysis Services translate MDX queries into SQL queries, which send the request to the Parallel Data Warehouse Control node for query processing.

Loading data using fastappend

When you need to complete a data load in a short time frame, and you don't need to rely on rollback, you can use fastappend. The fastappend command skips the temporary staging table and loads data directly into the destination table. In fastappend mode, the loader appends rows directly to the end of existing rows in the destination table without using a temporary table.

You can neither specify a staging database when using fastappend nor rely on rollback, so you must handle any recovery from a failed or aborted load as part of your own load process. However, fastappend speeds up loading, which is useful when you are under time constraints, such as having to complete a data load in two hours.

Using fastappend requires using the multi-transaction (`-m`) option.

Load data using fastappend:

1. Connect to a database using sqlcmd. The following example connects to a database called tpch1tb.
   ```
   sqlcmd -S "172.16.255.5,17001" -U sa -P password -I -d tpch1tb
   ```
2. Create a table named partsupp.

```
CREATE TABLE partsupp (
   ps_partkey      bigint,
   ps_suppkey      bigint,
   ps_availqty     integer,
   ps_supplycost   decimal(15,2),
   ps_comment      varchar(199))
WITH (Distribution = hash(ps_partkey));
```

The hashed distribution key, ps_partkey, distributes the data across all of the Compute VMs on the physical HSAxx servers.

```
C:\>sqlcmd -S "172.16.255.5,17001" -U sa -P [redacted] -I -d tpch1tb
1> ;
2> drop table partsupp;
3>
4> CREATE TABLE partsupp
5>   ( ps_partkey      bigint,
6>     ps_suppkey      bigint,
7>     ps_availqty     integer,
8>     ps_supplycost   decimal(15,2),
9>     ps_comment      varchar(199))
10> WITH (
11> Distribution = hash(ps_partkey));
12> ;
13> go
1>
```

FIGURE B1. SQLCmd

3. Execute dwloader on the servers that contain the source data.

```
dwloader -S 172.16.254.5 -U sa -P <password> -M fastappend -m -t "|" -r 0x7c0x0d0x0a -T
tpch1tb..partsupp -i d:\flatfiles.1TB\1\partsupp.tbl -R partsupp.out -rv 100 -b 100000
```

Parameters include:

- `-S` – IP address of the PDW Control VM.
- `-U` – User ID (in our example, sa).
- `-P` – System password.
- `-M fastappend` – Append rows directly to the end of existing rows in the destination table. Requires the multi-transaction (–m) option.
- `-m` – Commit loads in parallel. Required option for `-M fastappend`. Performs much faster than the default loading mode by bypassing the staging database. Has no rollback function, so any recovery from a failed or aborted load must be handled by your own load process.
- `-t "|"` – Delimit each field (column) in the row.
- `-r 0x7c0x0d0x0a` – Delimit row.
- `-T` – Target database.
- `-i` – Source data location (input).
- `-R` – File where you can find the rows which failed to load.
- `-rv` – Number or percentage of row rejections to allow before halting the load (in this example, 100).
- `-b` – Batch size; the number of rows to load per transaction. Defaults to 10,000.

The following example shows the beginning and end of `fastappend` command output.

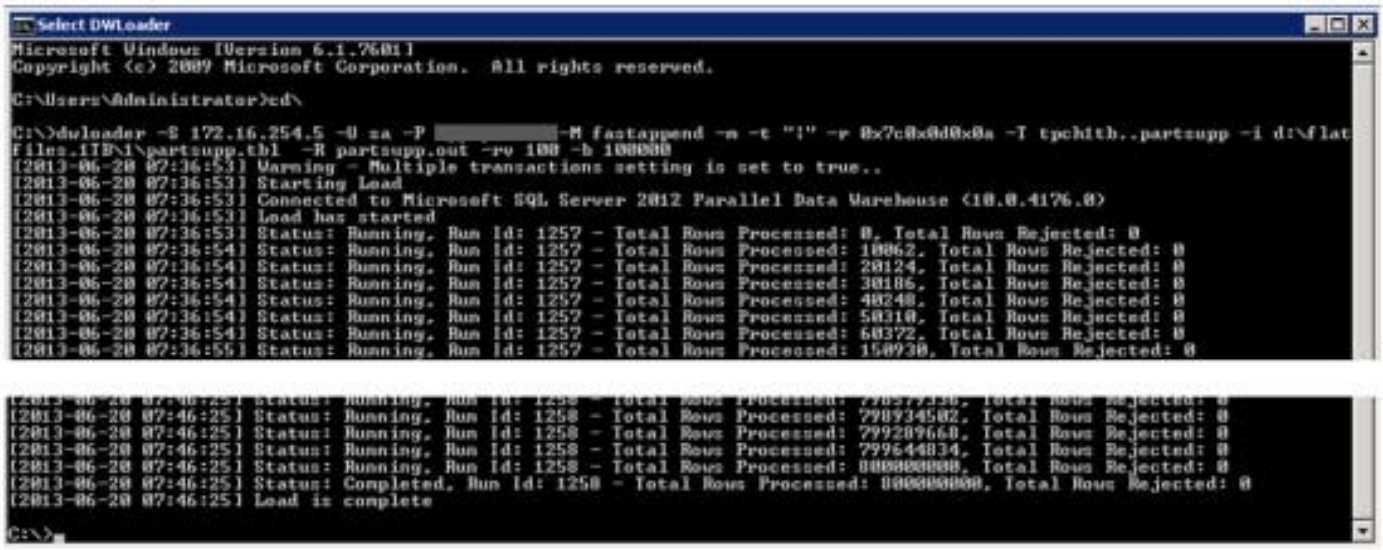

FIGURE B2. dWloader command-line loader

Adding Hyper-V nodes

Hyper-V Manager should list every server in your configuration. In a single-rack configuration, for example, you should see the following nodes:

```
<Fabric name>-HST01
<Fabric name>-HST02
<Fabric name>-HST03
<Fabric name>-HSA01
<Fabric name>-HSA02
<Fabric name>-HSA03
<Fabric name>-HSA04
```

```
<Fabric name>-HSA05
<Fabric name>-HSA06
<Fabric name>-HSA07
<Fabric name>-HSA08
```

Multiple-rack installations include a similar list for each rack. If Hyper-V Manager does not list all the nodes for your configuration, you can add them using the following procedure. The example below is to add missing node - HST04.

1. Log on to HST01 and launch Hyper-V Manager.
2. Right-click Hyper-V Manager in the left navigation, and then click Connect to Server.

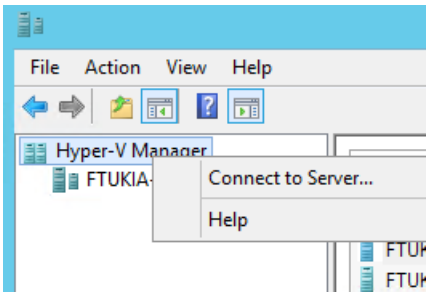

FIGURE B3. Connect to Server

3. Select *Another computer*, and then enter a server name into the text box.

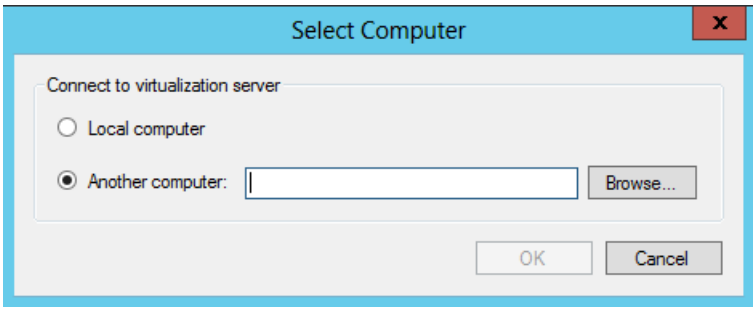

FIGURE B4. Select 'Another Computer'

You can also click Browse to locate the servers you want to add, and then click OK.

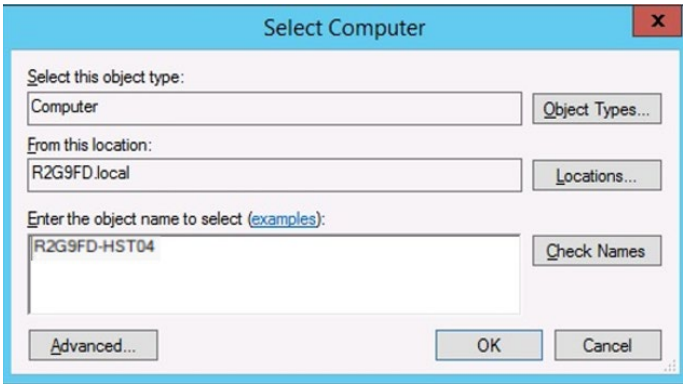

FIGURE B5. Browse for Computers

After connecting to all the servers, your Hyper-V Manager screen should appear similar to the following, depending upon the number of scale units you have purchased.

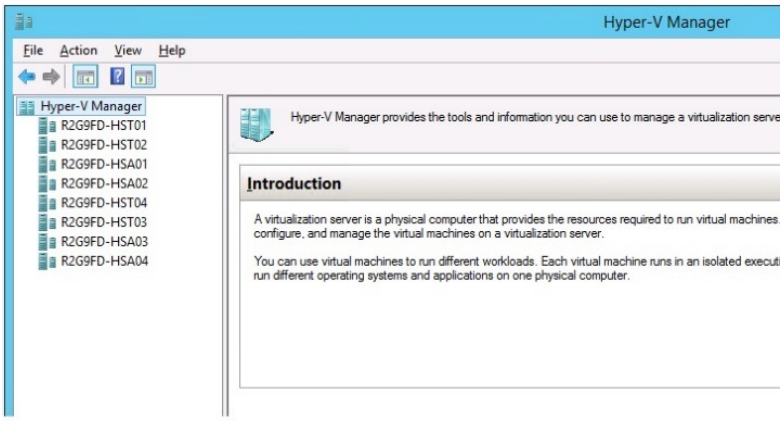

FIGURE B6. Hyper-V with the servers added

4. Verify that each HSA node running PDW has one compute node VM and one iSCSI VM:
   - HSA01 = CMP01 VM and iSCSI01
   - HSA02 = CMP02 VM and iSCSI02
   - etc.
5. Check the State column for each compute or data node VM and each iSCSI VM.
6. Start any VMs that are not shown as Running:
   a. Highlight the server in the left navigation pane.
   b. Right-click the VM in the center pane, and then click Start.

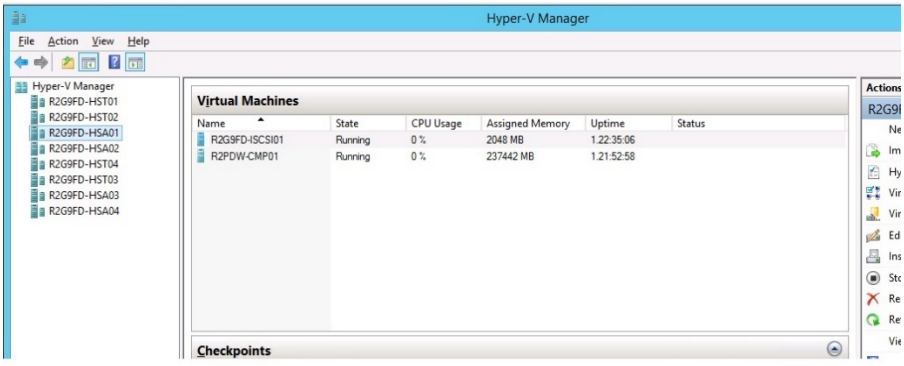

FIGURE B7. Hyper-V VM Alignment

7. If you find that any of the HSA servers are not aligned with the respective VMs (HSA01 with iSCSI01 and CMP01), proceed to next section to realign the VMs.

Hyper-V alignment procedures

Aligning servers and VMs

Under normal operation, HST01 runs three roles: the three clustered virtual machines, WDS, VMM, and CTL01. The remaining PDW HSTs exist for backup purposes in case of a system failure.

Occasionally, a VM becomes associated with an incorrect HST failover server, as shown below where CTL01 VM is running in HST02. CTL01 is a misaligned VM that you must move back to HST01 so that it is aligned with the appropriate server and exists with the three PDW VMs on HST01.

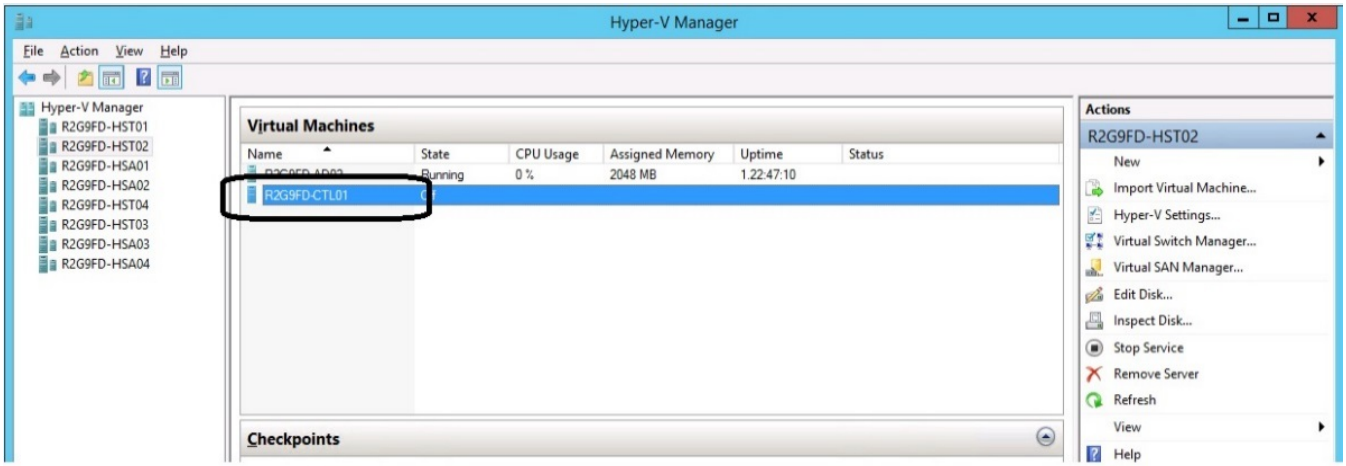

FIGURE B8. Misaligned VM

In a configuration where HST01 and HST02 are the only passive servers available, this misalignment can cause a reduction in availability levels because no passive servers are available if a compute node fails. This occurs because PDW resource requirements do not allow two separate compute VMs to exist on the same failover server.

Because this misalignment can cause a system failure, consider periodically verifying that all three VMs are properly aligned on HST01 and take steps to realign the VMs when necessary.

Moving a VM

To move VMs from one location to another:

1. Under Roles in the left navigation, right-click HST02.
2. Select Move → Quick Migration → Select Node.

NOTE

Live Migration is not supported and should not be used.

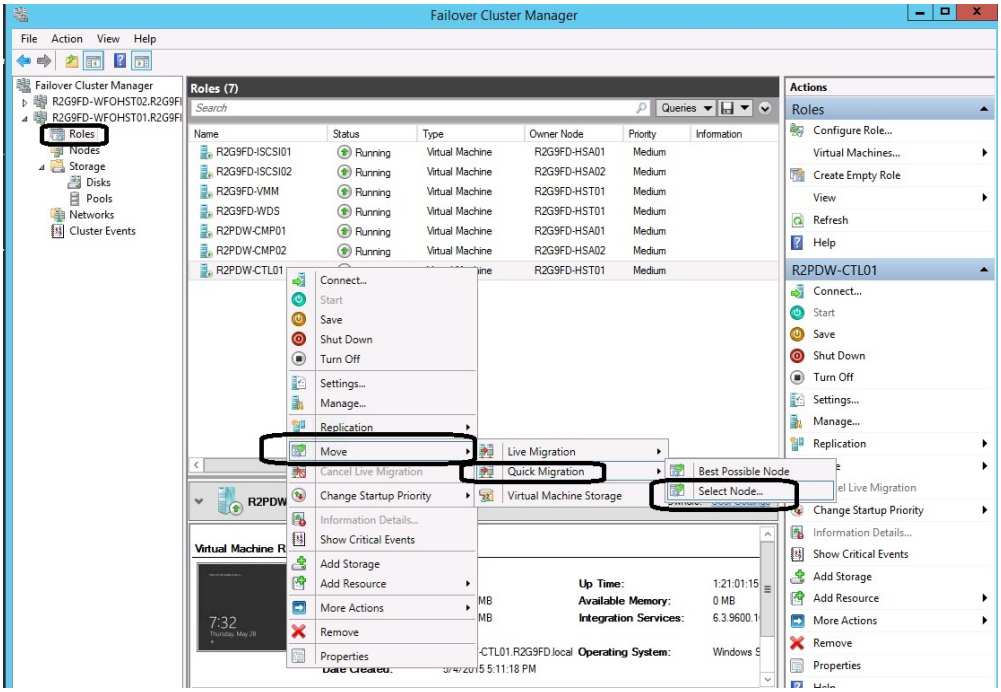

FIGURE B9. Using Quick Migration

In the Move Virtual Machine window, select HST01, and then click OK. Wait a few minutes for the operation to complete.

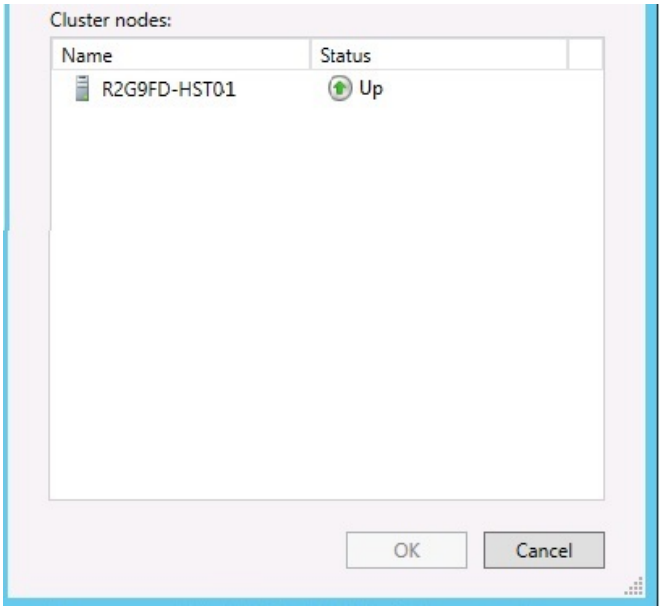

FIGURE B10. Selecting a target Hyper-V Host

Backing up and restoring SQL Server Parallel Data Warehouse data

Microsoft created backup and restore syntax for SQL Server Parallel Data Warehouse to capitalize on the parallel nature of the system. For more information about backup operations, see "Backup and Loading Hardware.

To provide for more flexible disaster recovery architectures, HPE ConvergedSystem 300 for Microsoft Analytics Platform enables you to connect multiple backup nodes. For more information see Backup and Restore.

Backup and restore syntax

Backing up the PDW databases on an HPE ConvergedSystem 300 for Microsoft Analytics Platform is a simple BACKUP DATABASE T-SQL command. As with all other query tools you would use, connect to the cluster IP address of the CTL01 VM, using port 17001. For example, if you used the sqlcmd command line utility, the following would be the command:

```
sqlcmd -S "172.16.252.5,17001" -U sa -P <yoursapassword> -I -d AdventureWorksPDW2016
```

Parameters included in the example are defined as follows:

`-S` – IP address of the Control node CLUSTER that the Admin Console, operation commands, and users use for access

`-U` – User ID (in the example, sa)

`-P` – System password

`-I` – Enabled Quoted Identifiers

`-d` – Name of the database to back up

Unlike a standard SQL Server backup, a SQL Server PDW backup will be a collection of files. Therefore, when running the backup command, you will specify the UNC path to a backup directory. It is strongly recommended that this IP address be of a server attached over the InfiniBand network as described earlier in this document. If so, then your high-speed backup will flow onto the backup server over the InfiniBand connection. This will be on a separate subnet from your Ethernet connections.

In our example, if we had a backup/data loading server, it could have two InfiniBand (IB) connections to the HPE ConvergedSystem 300 for Microsoft Analytics Platform, one on the IB1 network (172.16.254.99), and one on the IB2 network for redundancy (172.16.255.99). To run the backup, we'd have a file share (for example, PDWBackups), and the command would look like:

```
BACKUP DATABASE AdventureWorksPDW2016 to DISK='\\172.16.254.99\PDWBackups\AdventureWorksPDW2016';
```

Restoring a database is a reverse of the same process. The exception is that if you need to restore the master database, you must use the dwconfig utility.

To restore the AdventureWorksPDW2016 database we backed up above, you would run:

```
RESTORE DATABASE AdventureWorksPDW2016 FROM DISK='\\172.16.254.99\PDWBackups\AdventureWorksPDW2016';
```

For additional details about backup and restore syntax, refer to Backup and Restore database (Parallel Data Warehouse).

Moving backup data to a corporate-wide backup system

Copy the backup directory files to a corporate-wide backup or tape library for disaster recovery purposes.

If you follow the advice in this guide and use an InfiniBand connected backup server, moving a backup set off the backup node can negatively impact restore performance. Whenever possible, perform restore operations from the backup files as they were originally stored on the InfiniBand-connected server. If possible, copy the backup directory from your corporate-wide backup or tape library to the recommended server, and then perform the restore for optimal performance.

Adding users to the appliance

Add new SQL Server 2016 Parallel Data Warehouse logins using the CREATE LOGIN SQL statement. For more information, see SQL Reference, Security Statements, CREATE LOGIN.

As of Appliance Update 1, SQL Server PDW supports both Windows Integrated security logins and SQL Server security logins. Please follow the guidance from the APS guide in creating your security logins. Do note that the admin console supports both SQL Server security and Windows Integrated Security userid/passwords, but does require provisioning and permissions to run the admin console.

Device passwords

During installation, the Hewlett Packard Enterprise support team will reset all passwords according to site requirements. Use the following procedures if you need to change the passwords for the hardware or HPE software components of the HPE ConvergedSystem 300 for Microsoft Analytics Platform.

NOTE

Hewlett Packard Enterprise strongly recommends using the same password throughout the HPE ConvergedSystem 300 for Microsoft Analytics Platform.

Changing passwords requires the IP addresses of the various system components, including iPDUs, iLO, and InfiniBand switches.

Finding device IP addresses

To find the device IP addresses:

1. Log on to HST01 (physical server, not the VM).
2. Open the C:\PDWINST\Media folder.
3. Locate and open `HPEApplianceDetails.xml` in a text editor. The file lists the IP addresses for the system components, as shown in the following example:

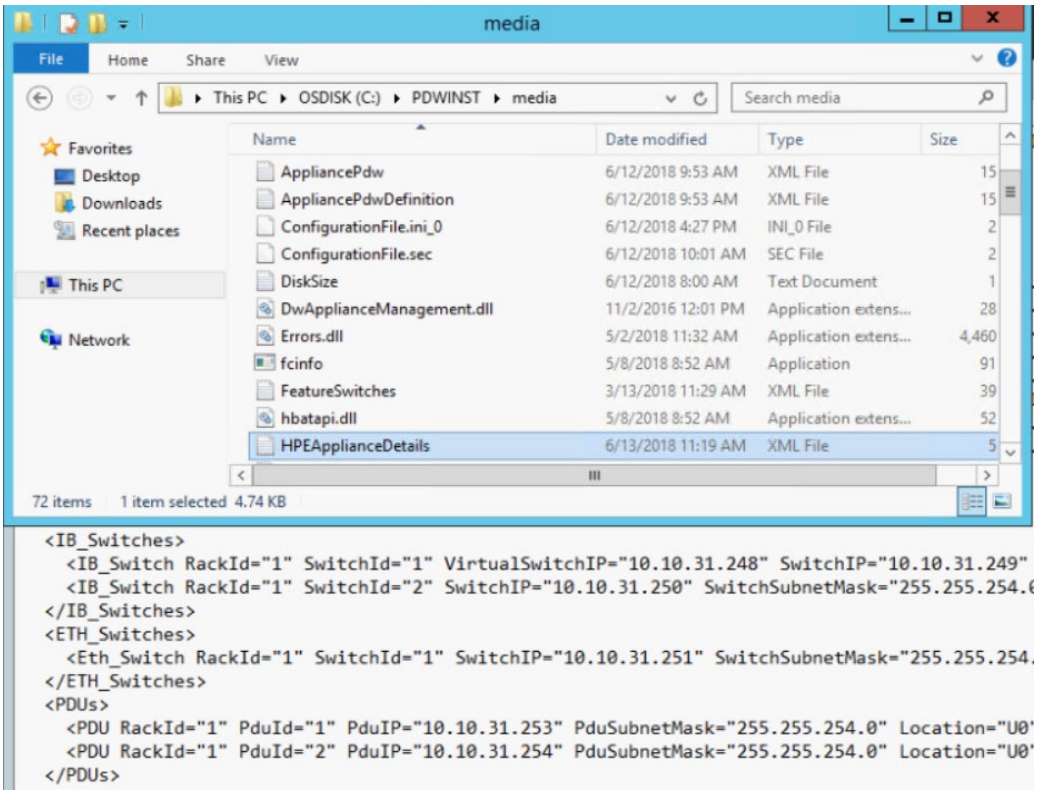

FIGURE B11. Finding HPEApplianceDetails.xml

iLO passwords

To change the iLO password you need to know the iLO IP addresses. You can find all server iLO IP addresses on HST01 in the `BmcAddress` file under `C:\PDWINST\Media\HPEApplianceDetails.xml`.

1. Open a browser window.

2. Enter one iLO IP address into the browser, and then press Enter.

The browser might return an error regarding a security certificate for this device. To avoid this error, replace the default certificate with a trusted certificate. For more information, see the Appliance Configuration (Analytics Platform System)

3. Click "Continue to website (not recommended)" to launch the iLO HPE ProLiant login screen.
4. Log in to the iLO 5 HPE ProLiant as Administrator using the current password.

In the Information - iLO Overview screen, select **Administration** to view the default **User Administration** menu.

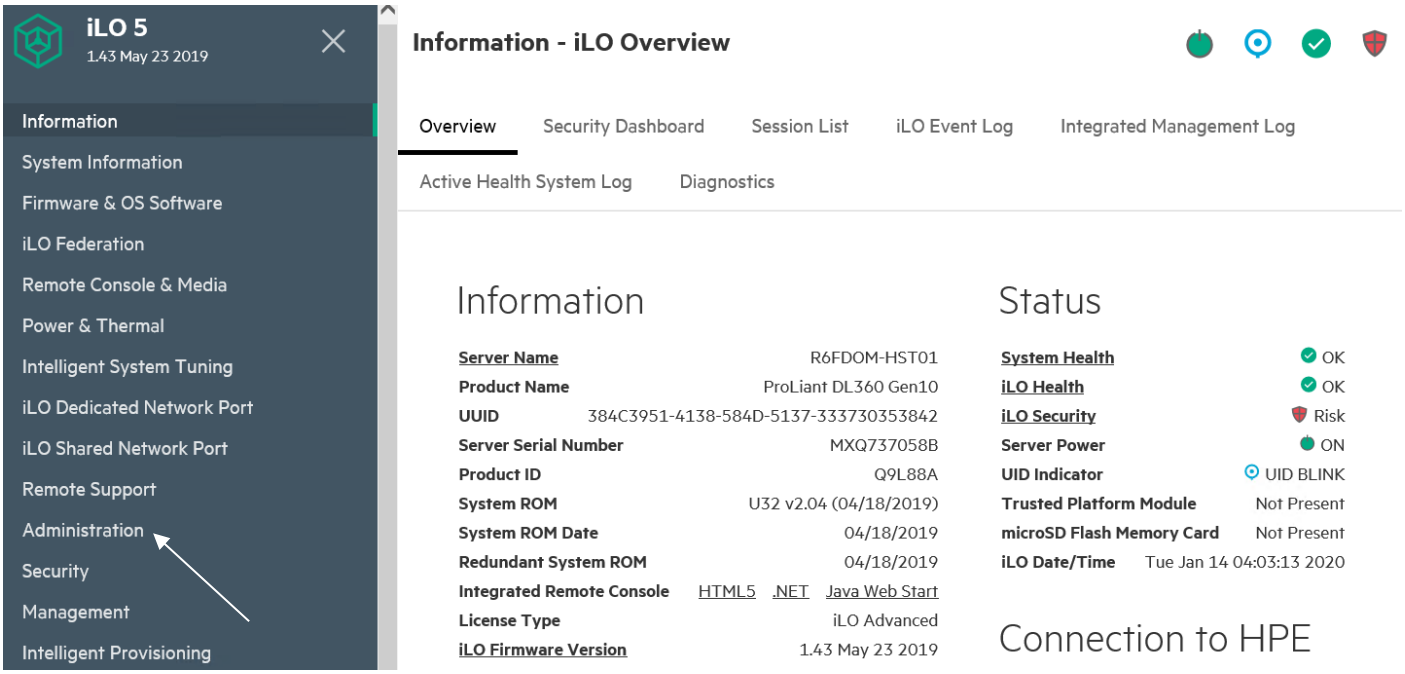

FIGURE B12. iLO User Administration

5. Under Local Users, select the checkbox next to Administrator, and then click **Edit**.

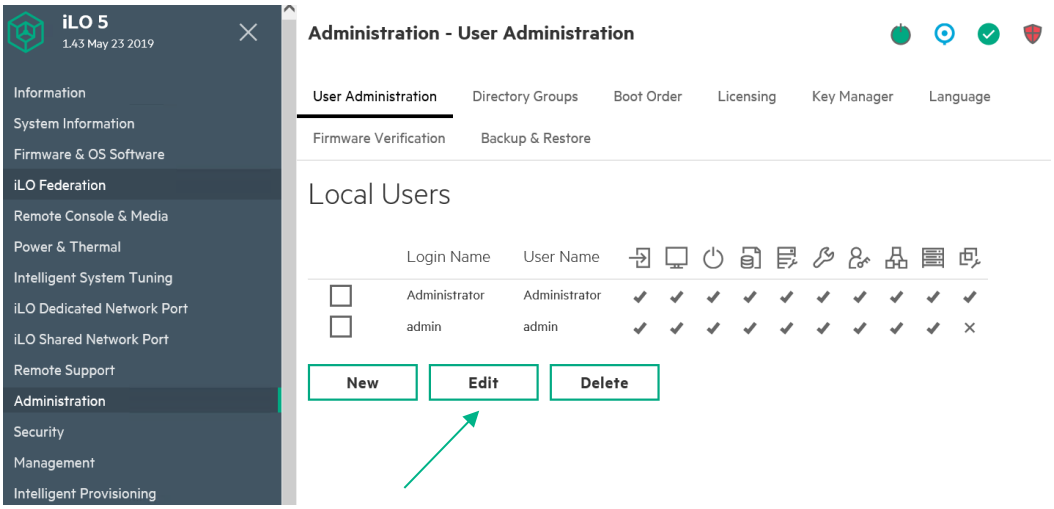

FIGURE B13. iLO Administrator Account Edit

6. Click on the checkbox for "Change password"

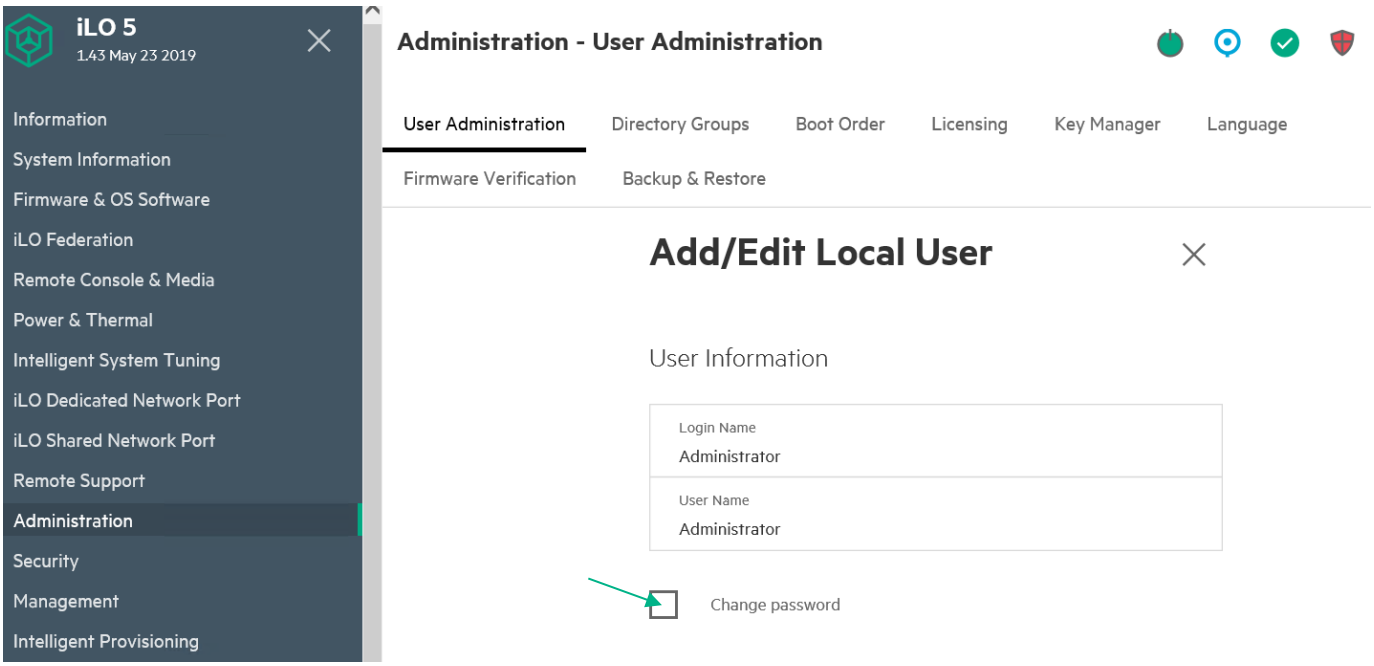

FIGURE B14. Edit Administrator's password

7. Enter the new password information, and then click Update User.

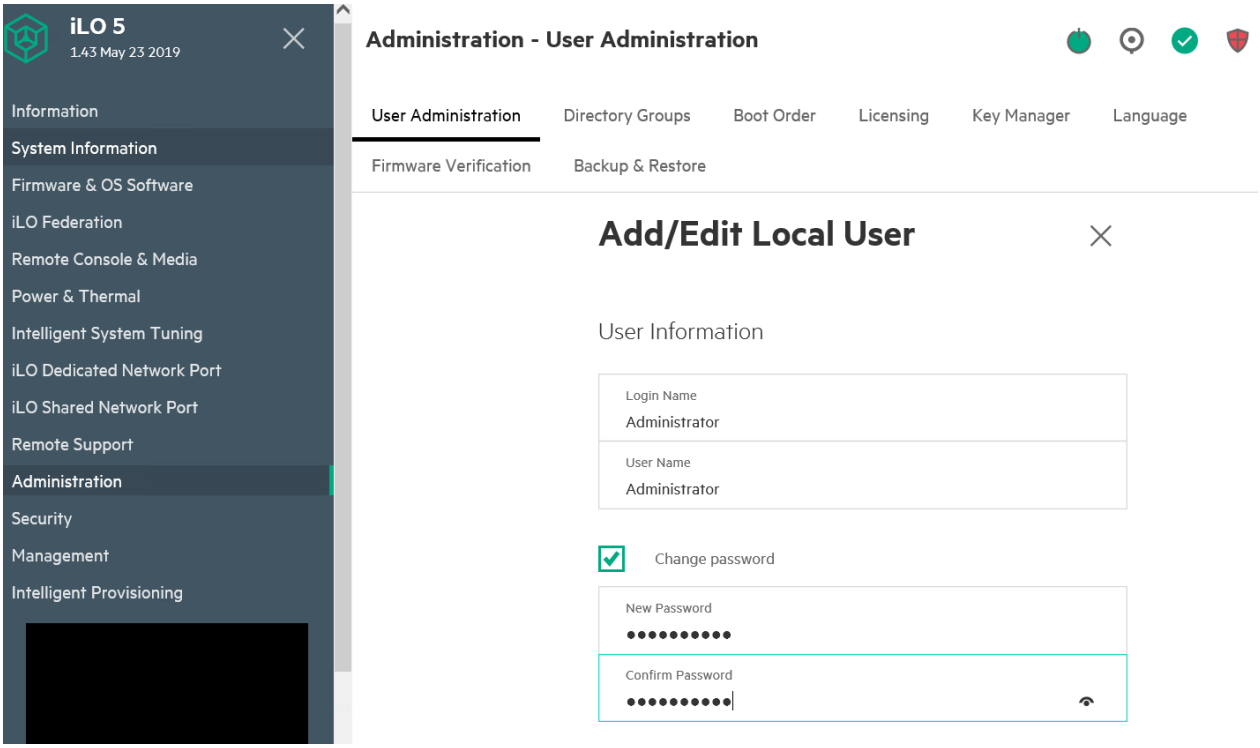

FIGURE B15. iLO password save

8. Repeat this procedure for the iLO on each physical server in the system.

Ethernet switch passwords

Enabling Ethernet switch browser access

Requirements at your site might include disabling browser access to the Ethernet switches for security reasons. You can enable browser access to simplify any necessary password changes. Disabling browser access after you complete the password changes is optional.

1. Log on to the Ethernet switch as admin.
2. Enter *system-view* at the system prompt, and then press Enter.
3. Enter *ip http enable*, and then press Enter.
4. Enter *save force*, and then press Enter.

NOTE

Wait while the system validates the file. Continue with the next step when you see "Configuration is saved to device successfully."

---

5. Enter *local-user admin*, and then press Enter.
6. Enter service-type web at the luseradmin prompt, and then press Enter.
7. Enter *save force*, and then press Enter.
8. Wait while the system validates the file.
9. Continue when you see "Configuration is saved to device successfully."

Verify that the change was successful.

10. Enter *display current-configuration* at the system prompt, and then press Enter.
11. Scroll through the data returned by the command to local-user admin.
12. Verify that service-type web appears in the list as shown below.

```
#
local-user admin
 password cipher $c$3$lNkvR8QxIYQAqdTqDdoewEMVlONrXQqFt8XpDe4p  authorization-attribute level 3
service-type lan-access  service-type ssh terminal  service-type portal  service-type web
```

Ethernet switch passwords

To change the Ethernet switch password, you need to know the Ethernet switch IP addresses. To find them, see Finding device IP addresses. You might also need to enable browser access to the Ethernet switch (see Enabling Ethernet switch browser access).

To change the Ethernet switch password using PuTTY as the Telnet interface (or something similar):

1. Open a PuTTY session.
2. Enter one Ethernet switch IP address into the session, and then press Enter.
3. Then log in to the switch using the admin account and your current credentials:

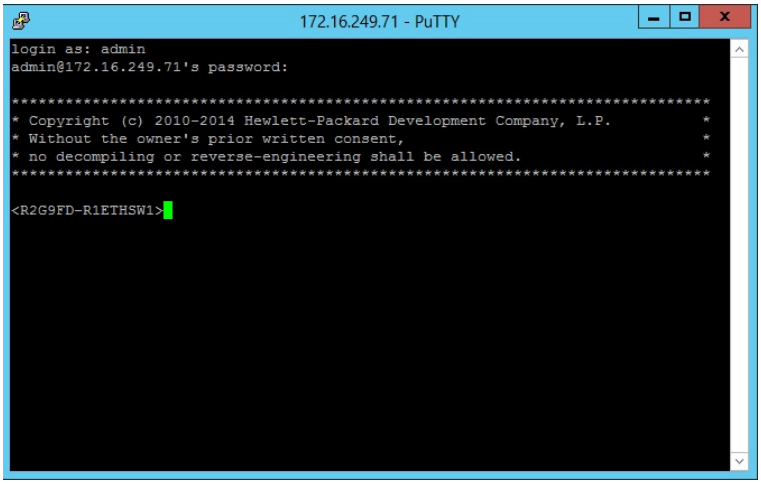

FIGURE B16. Switch login

4. Switch to the System view by typing in “sys” and then press Return.
5. Next type in “local-user admin” to get into the interface for changing attributes for the admin account:

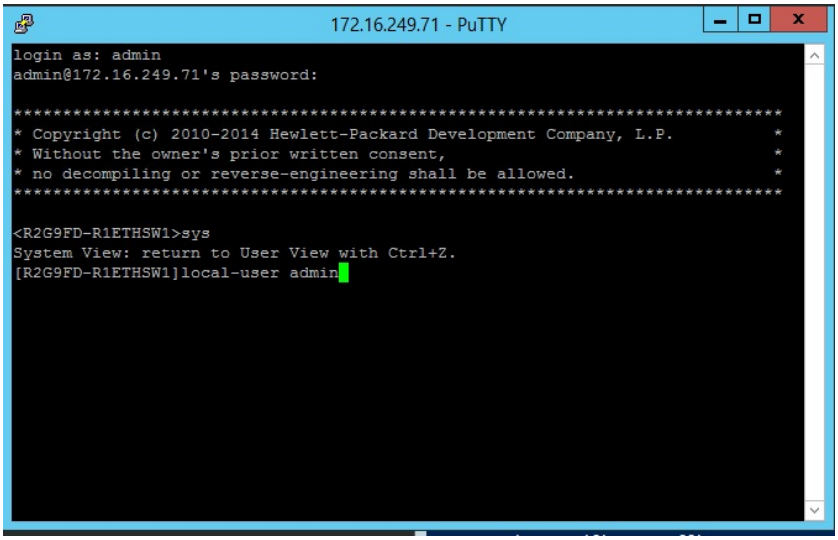

FIGURE B17. Admin user context

6. Next type in “password” to get to the password dialog. Then simply type the new password, then press Return. You will be prompted to confirm. After re-typing the password, it should be successfully changed:

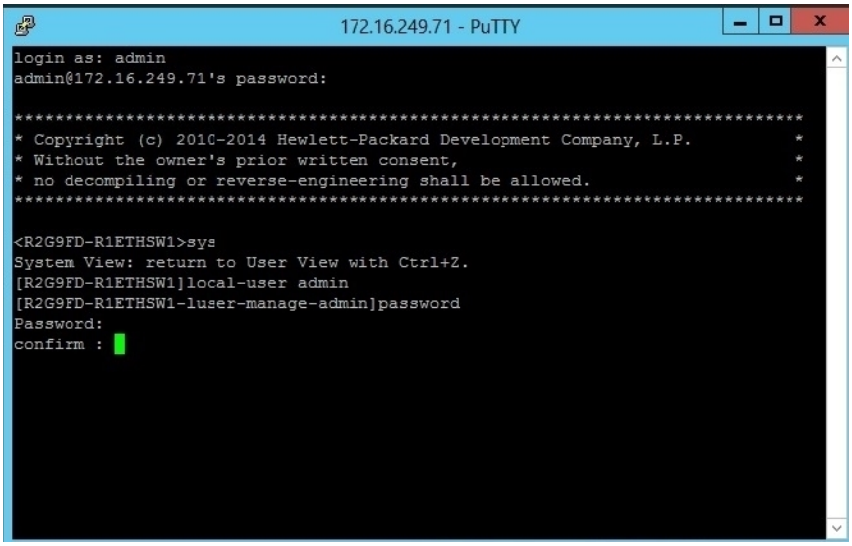

FIGURE B18. Password Change

EDR InfiniBand switch passwords

Using PuTTY to change EDR InfiniBand switch passwords

To change the EDR InfiniBand passwords you need to know the EDR InfiniBand IP addresses. To find them, see Finding device IP addresses.

If you have disabled web access on the InfiniBand switches for security reasons, you can use a serial cable to connect a laptop (or use one of the HST servers) to the switch.

If you are able to use a browser, see Using a browser to change EDR InfiniBand switch passwords.

To change the InfiniBand password using PuTTY as the Telnet interface:

1. Connect your laptop to the InfiniBand switch, and then use PuTTY to access the switch.
2. Log on as admin using your current password.
3. At the command prompt, enter *enable*, and then press Enter.
4. Enter *configure terminal*, and then press Enter.

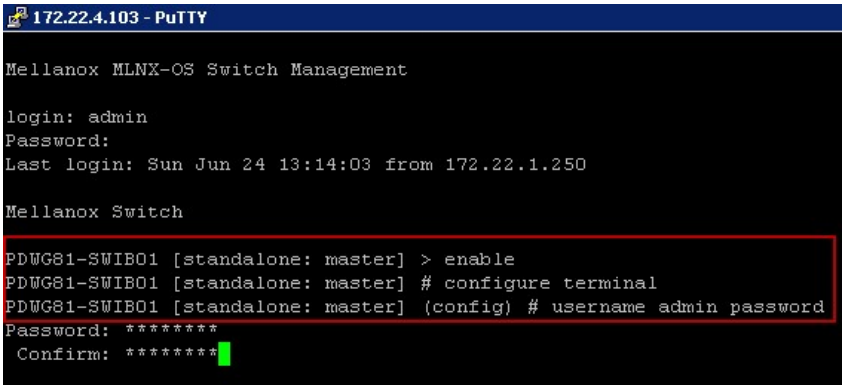

FIGURE B19. InfiniBand Switch Login

5. Enter the new password at the Password prompt, and then press Enter.

6. Enter the new password again at the Confirm prompt, and then press Enter.
7. Enter *exit*, and then press Enter.
8. Enter *exit*, and then press Enter.
9. Repeat this procedure for each EDR InfiniBand switch in the configuration.

Using a browser to change EDR InfiniBand switch passwords

To change the EDR InfiniBand password you need to know the EDR InfiniBand IP addresses. To find them, see Finding device IP addresses.

If you are unable to use a browser, see Using PuTTY to change EDR InfiniBand switch passwords.

Change the System Administrator password using a browser:

1. Open a browser window.
2. Enter one InfiniBand IP address into the browser, and then press Enter.
3. The browser might return an error regarding a security certificate for this device. To avoid this error, replace the default certificate with a trusted certificate. For more information, see Appliance Configuration Tasks (SQL Server PDW
4. Click "Continue to website (not recommended)" to launch the Mellanox MLNX-OS Management Console.
5. Log in to the switch as admin using the current password.
6. From the Summary screen, select the Security tab.

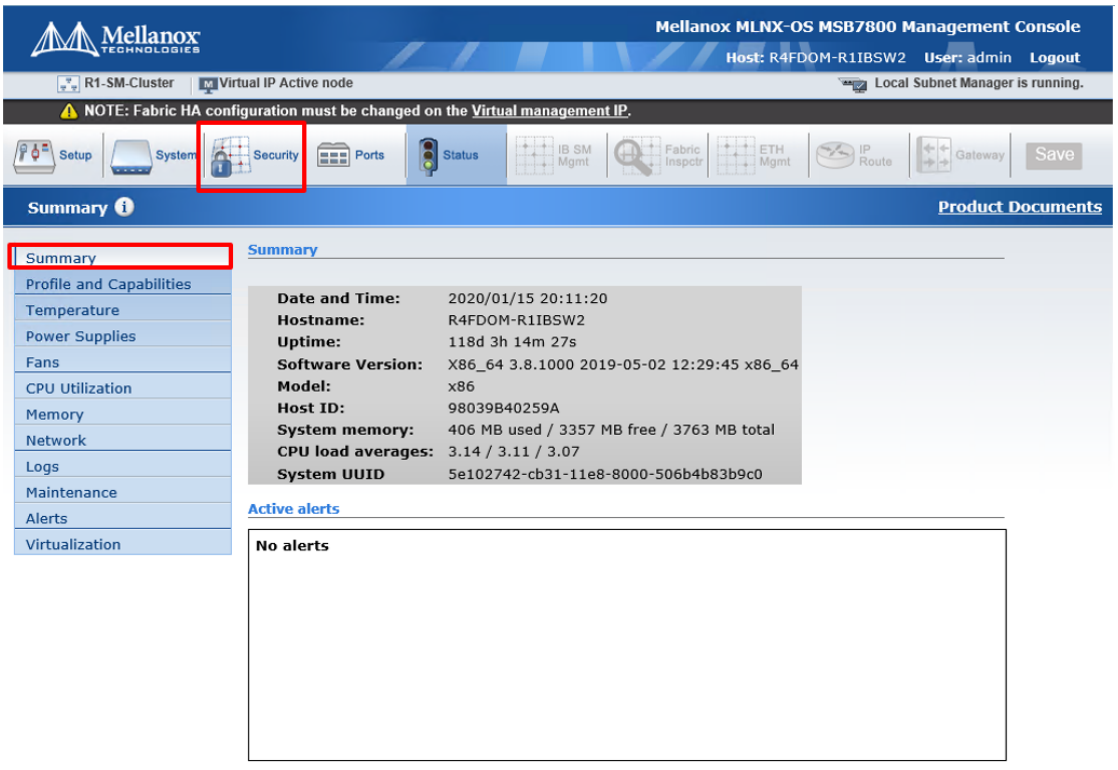

FIGURE B20. InfiniBand Security Management

7. Click Edit next to the admin username under User Accounts.

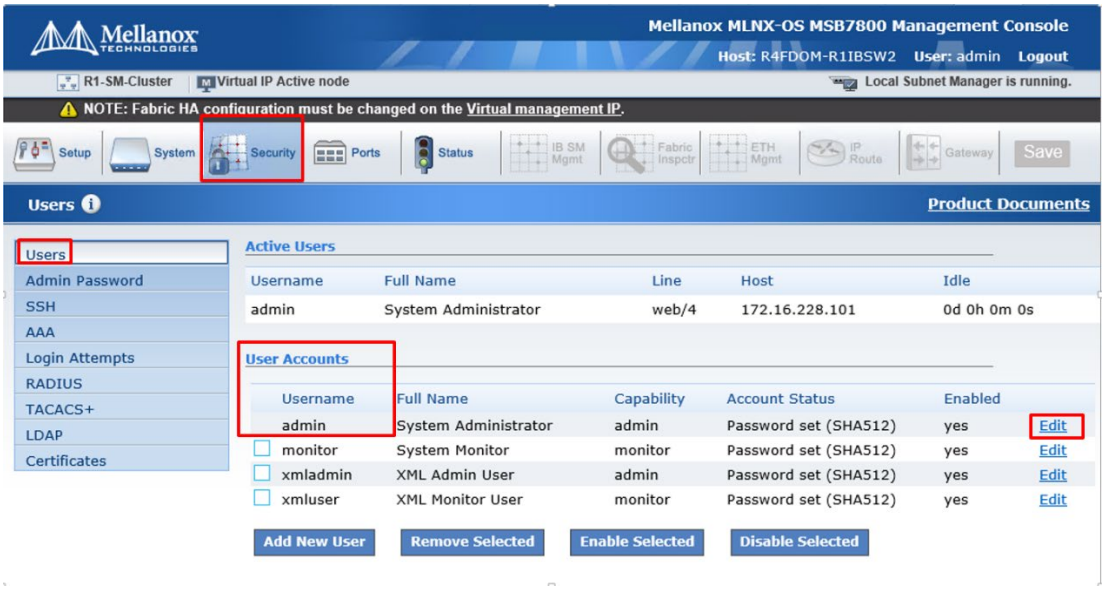

FIGURE B21. InfiniBand Switch User Management

8. Enter the new password in the **New password field**, and then enter the same password into the **Confirm new password** field.
9. Click **Apply**, and then click **OK**.

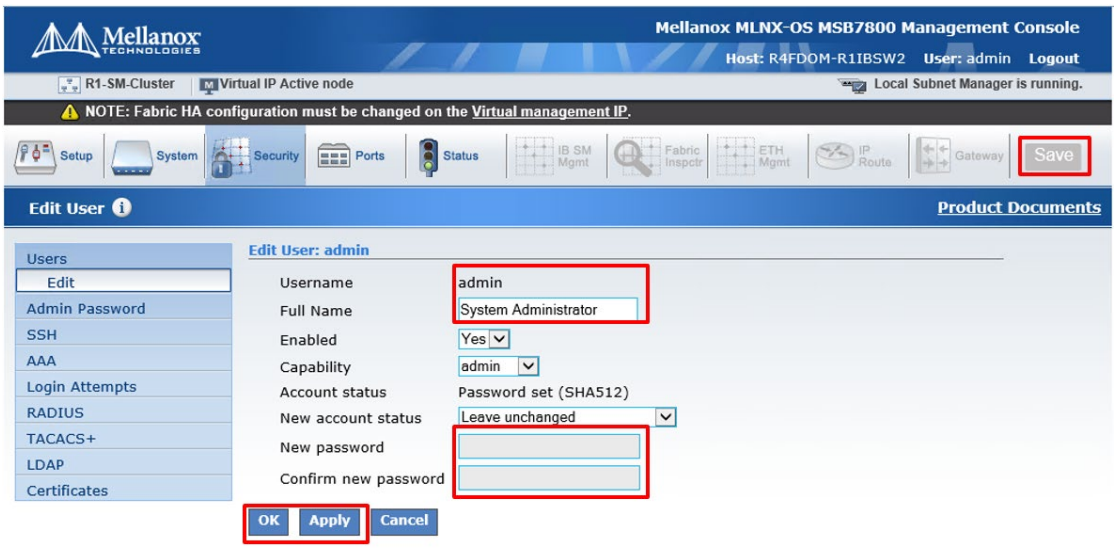

FIGURE B22. InfiniBand Switch Password Change

10. Click **Save** in the upper right corner.
11. Change the System Monitor password:
12. Click **Edit** on the System Monitor line under User Accounts.

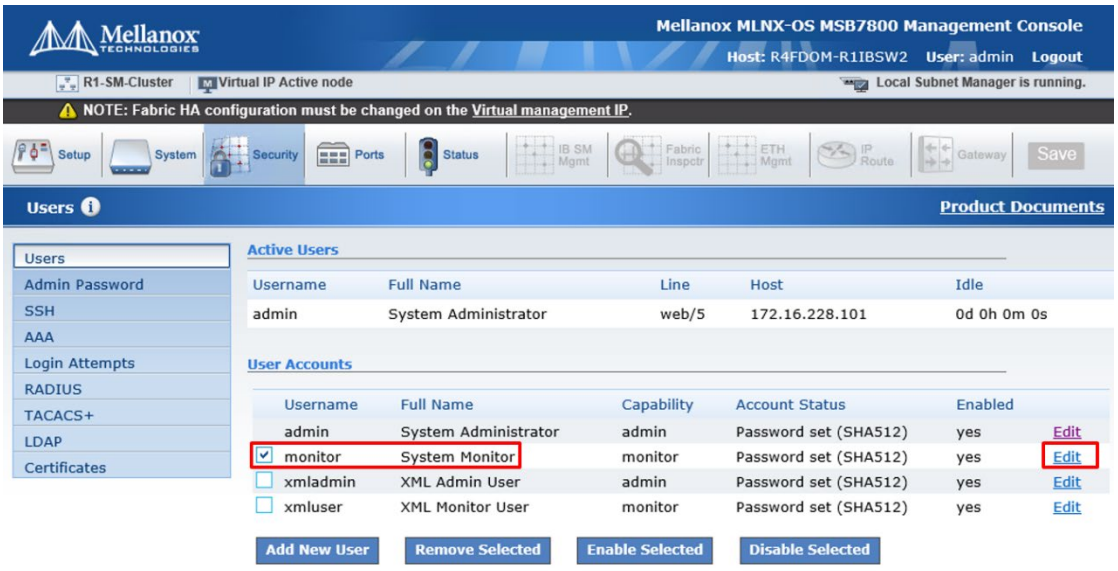

FIGURE B23. Edit System Monitor User

13. Enter the new password in the **New password** field, and then enter the same password into the **Confirm new password** field.
14. Click **Apply**, and then click OK.

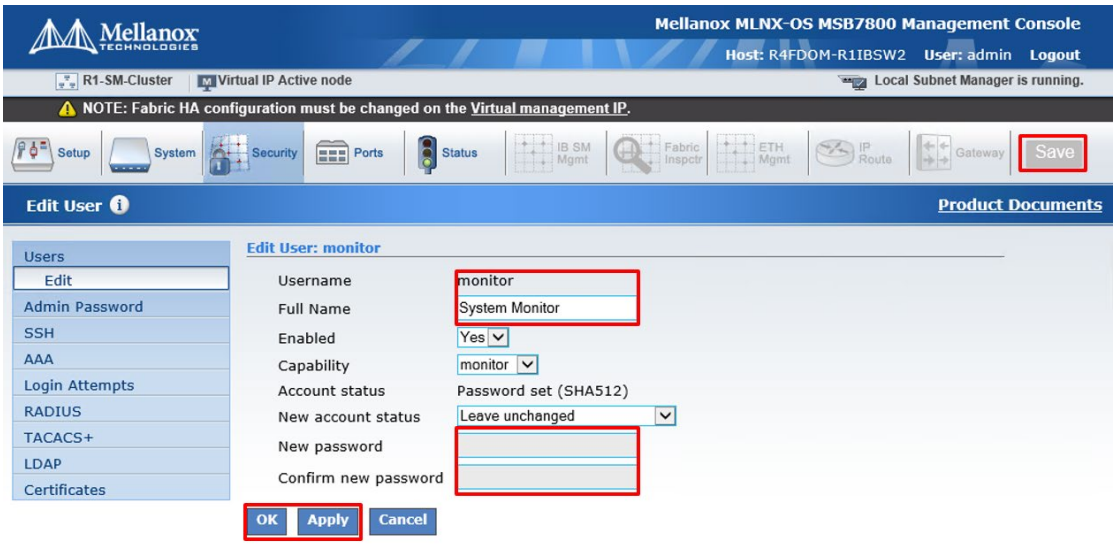

FIGURE B24. Changing the System Monitor Password

15. Click **Save** in the upper right corner.
16. Click **Logout** in the upper right corner of the MLNX-OS Management Console to log off of the switch, and then log on again to verify the new passwords.

Repeat this procedure for each EDR InfiniBand switch in the configuration.

Metered PDU IP & passwords

Using PuTTY to change IP & PDU passwords

For HPE Metered PDUs do the following steps to set the IP address and to change the admin password.

NOTE
The serial+RS485-1 port takes a special DB9 pin configuration. Ensure your RJ-DB9 cable matches the pinout assignments per below. You may use an RJ45-DB9 adapter that matches the pin-outs below.

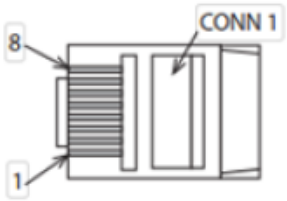

CONN 1		CONN 2
Pin	Description	Pin
1		1
2	RS232 RX	2
3	RS232 TX	3
4		4
5	Ground	5
6		6
7		7
8		8
		9

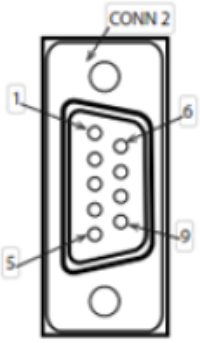

FIGURE B25. Serial+RS485-1 DB9 pin configuration

1. Connect an RJ45-DB9 cable between the port labeled “Serial+RS485-1” on the Network Management Module of the PDU and the serial connector on a host computer. This connection will be used to configure and access the PDU through a terminal emulation program, as well as upgrade the PDU.
2. On the host computer start an emulation program like Putty. Select the serial connector on the host computer to which the RJ45-DB9 cable is attached. Use the following connection parameter values.
3. Gen10 servers default serial port is COM2

   Bits per sec – 115200
   Data bits – 8
   Parity – None
   Stop Bits – 1
   Flow Control – None

4. Use the default initial login indicated below. The username and password are both case sensitive.
   - Username: admin
   - Password: <APS default password from factory>
5. The “HPE” prompt appears after you have logged in, ready to enter the CLI command.
6. Configure Network IP:
   - Type *net tcpip static <ip address> <netmask> <gateway>*, then press Enter
   - Enter N to skip reboot
7. Configure web access:
   - Type *net https on*, then press Enter

- Enter **N** to skip reboot

8. Configure telnet access:
   - Type *net ssh off*
   - Type *net FTPs off*
   - Enter **Y** to confirm and the PDU Mgmt module will reboot and save all new settings.
9. Change the ***admin*** account password:
   - Using a browser on the host computer, log into the PDU Management GUI.
   - Click on the ***admin*** user to display the drop-down menu and select **Change Password**.

FIGURE B26. Admin drop-down menu

   - Enter Current Password: <APS default password from factory>
   - Enter new password: xxxxxxxxx
   - Enter confirm new Password: xxxxxxxxx
   - Click on **Change Password** Icon to apply new password
   - Click the ***admin*** user again to display the drop-down menu and select **Log Out**
10. Move to the next PDU to configure until all PDUs have been configured.

RESOURCES AND ADDITIONAL LINKS

HPE Reference Architectures, hpe.com/info/ra

HPE Servers, hpe.com/servers

HPE Storage, hpe.com/storage

HPE Networking, hpe.com/networking

HPE Technology Consulting Services, hpe.com/us/en/services/consulting.html

HPE ConvergedSystem 300 for Microsoft Analytics Platform (includes QuickSpecs), https://www.hpe.com/products/cs300aps

Microsoft Analytics Platform System Product Overview, https://www.microsoft.com/aps

Online Documentation, https://docs.microsoft.com/en-us/sql/analytics-platform-system/index?view=aps-pdw-2016-au7

To help us improve our documents, please provide feedback at hpe.com/contact/feedback.

4aa5-9215enw , version 8.0, January 2021

Hewlett Packard Enterprise