Reference Architecture

HPE Reference Configuration for Elastic Platform for Analytics (EPA)

Modular building blocks of compute and storage optimized for modern workloads

Contents

Executive summary

As organizations strive to identify and realize the value in Big Data, many now seek more agile and capable analytic systems to deliver on that value. Enterprises are looking to accelerate their analytics workloads to drive real-time use cases such as fraud detection, recommendation engines, and advertising analytics. The separation of processing from resource management with YARN has driven new architecture frameworks for Big Data processing, such as Lambda, Kappa and SMACK (Spark, Mesos, Akka, Cassandra and Kafka). These frameworks are designed to extend Hadoop beyond the often I/O intensive, high latency MapReduce batch analytics workloads, to also address real-time and interactive analytics. Technologies like Spark, NoSQL, and Kafka are critical components of these new frameworks to unify batch, interactive and real time Big Data processing. These technologies typically have different capacity and performance requirements to scale out processing, and need a variety of options for compute, storage, memory and networking. This presents a significant challenge to the traditional framework in Hadoop; where collocated compute and storage inhibit scale of either resource independently, and multi-tenancy of workloads creates scheduling and performance challenges.

With the increase in data generated outside of the data center, or at the edge, businesses are looking at how to capture, analyze and grow new services from that data. The challenge lies in harnessing the data and processing it, while streaming for real-time or near-real-time analysis and response. An end-to-end data pipeline that extends from the edge to the data center needs to exist in order to accomplish this.

The HPE Elastic Platform for Analytics (EPA) is designed as a modular infrastructure foundation to address the need for a scalable multi-tenant platform, by enabling independent scaling of compute and storage through infrastructure building blocks that are workload optimized.

Document purpose: This document describes different deployment options for the HPE Elastic Platform for Analytics platform (previously referred to as the HPE Big Data Reference Architecture or BDRA) using modular building blocks. This white paper also provides suggested configurations that highlight the benefits of a building block approach to address the diverse processing and storage requirements typical of modern analytics environments. This version of the document has been updated to include HPE ProLiant Gen10 servers, HPE Apollo Gen10 servers and updated networking building blocks.

Target audience: The intended audience of this document includes, but is not limited to IT managers, pre-sales engineers, services consultants, partner engineers and customers that are interested in deploying the various software components of the larger Hadoop ecosystem with infrastructure solutions from Hewlett Packard Enterprise.

This white paper describes a project developed by Hewlett Packard Enterprise in November 2018.

Introduction

Since being founded in 2006, Hadoop has proven to be cost-effective at solving the challenge of managing growing data across organizations. Hadoop provided a batch processing framework (MapReduce) and an efficient, software-defined-storage file system (HDFS) for storing data whose velocity, volume, and variety has characterized it as “big”. Rather than buying expensive “latest and greatest” machines that each performed a large workload, distributing the data and sharing the processing across a cluster mitigated large purchases of super-computers for these demanding tasks. Hadoop contributors favored data locality, with collocated compute and storage at the node. Each server added more compute and storage capacity, scaling linearly. This is what many consider to be a traditional, symmetric architecture where each server is configured identically.

Hadoop 2.0 extended the batch processing environment to improve upon the in-place analytics engines, as well as provide additional support for analytics on streaming workloads. It also introduced the concept of multi-tenancy for these varied, disparate workloads through a new resource manager (YARN). YARN’s resource management framework using containers allows for more granular allocation of compute and memory, and targeting workloads to specific groups of servers through the use of YARN labels, to address the diverse computing requirements of newer compute and memory intensive workloads.

Many organizations shared a vision of having all workloads, analytics and applications running on a common data set, on a scalable multi-tenant platform for all data and analytical workloads. This exercise typically starts with consolidating data using Hadoop as a data repository and for simple workloads like ETL and pre-processing data. The rate of change and variety in the Hadoop software stack to accelerate a diverse set of analytics use cases beyond just ETL and ELT has driven organizations to rapidly evaluate and deploy these technologies to build new business capabilities. They want enterprise-grade performance to run diverse set of workloads, consolidate data and infrastructure with the ability to scale these workloads across a common, elastic, shared and flexible infrastructure.

Several trends in data analytics are playing a role in the solutions HPE is delivering today and in the future. Some are adoption areas from our customers, some are technologies that are emerging and are being more effectively utilized in these big data analytics solutions. Among them:

- Consumption-based experience – A managed infrastructure with flexible growth and pricing structures
- Stream analytics – Moving beyond MapReduce to analyze data in flight from the edge, to the core and into the cloud
- AI/Deep learning – Machine learning algorithms are being applied to big data to solve business problems in every industry
- Multi-tenancy – Drive toward reducing cluster sprawl with density-optimized servers serving diverse workloads on a single data platform
- Hyperscale – Faster networks and flash storage allow for independent scaling of compute and storage, enabling new platforms for analytics
- Containers – Enable infrastructure and business agility with Docker, Mesos and Kubernetes

Based on these trends an end-to-end data pipeline must be architected to deal with these obstacles. Legacy data and analytics systems are ill-equipped to handle the new types of Big Data workloads. Existing data lakes and batch-oriented Big Data clusters are limited architecturally. Siloed Hadoop clusters have created "ponds" of data, leaving a fragmented data landscape. A majority of customers cannot fully exploit their data nor can they analyze it fast enough. There is an overwhelming flow of new data types coming from the edge, particularly with AI. Data protection and governance is still evolving, especially in a hybrid world, and is becoming more important as enterprises comprehend the value of their data.

Figure 1 illustrates an example of the data flow from edge to core to cloud, routed through a data pipeline that provides an infrastructure for data to not just flow bi-directionally but also to allow for the implementation of analytic processes in real-time, near real-time, at rest and also AI modelling.

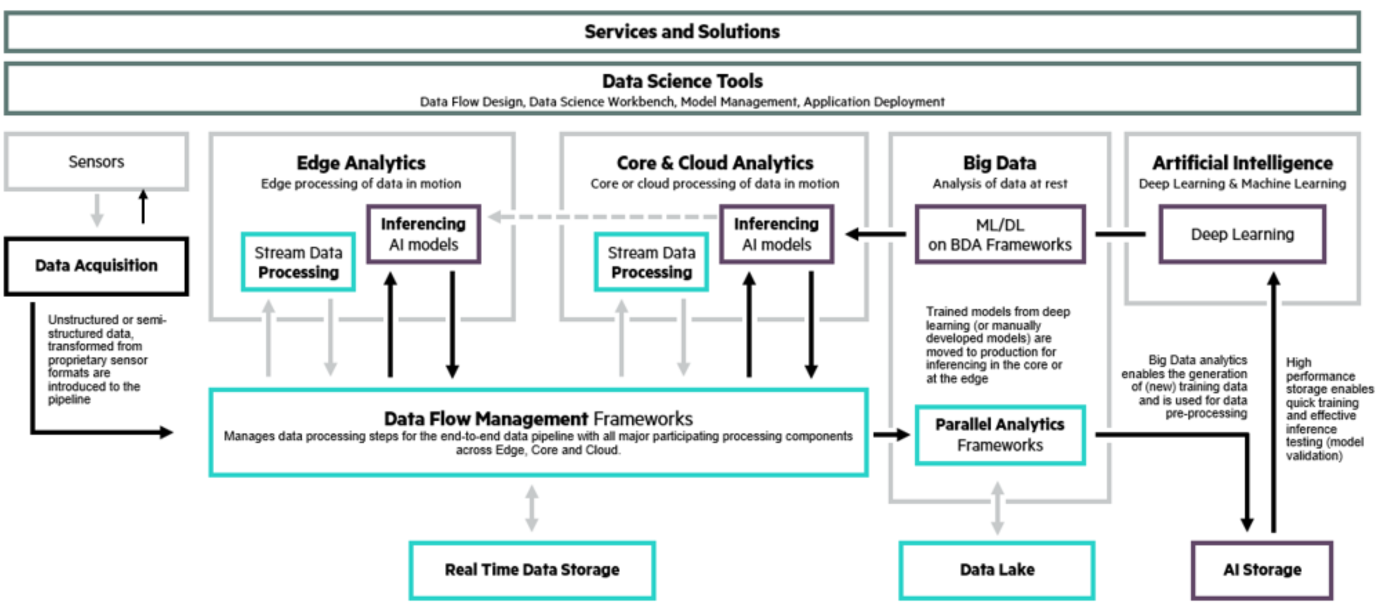

Figure 1. HPE Elastic Platform for Big Data Analytics (EPA) end-to-end data pipeline

As Hadoop has evolved, many customers have experienced trapped resources as workload needs have changed. HPE has recognized the challenges and limitations of deploying Hadoop in the traditional manner, and has introduced the HPE Elastic Platform for Analytics to address some of these challenges. HPE EPA provides a modular platform foundation for deploying analytics frameworks and Hadoop workloads to build end-to-end data analytics pipelines. This solution solves the common challenges many customers face through a robust, yet flexible, offering that enables organizations to maximize the performance, infrastructure and analytics abilities of Hadoop on an enterprise-grade HPE solution that can scale with their evolving business needs.

HPE Elastic Platform for Analytics (EPA)

HPE's Elastic Platform for Analytics offers an asymmetric architecture and a different approach to building a data pipeline optimized for Big Data Analytics. The asymmetric architecture consists of separately definable and scalable compute and storage resources, connected via HPE high speed network components (25Mb, 40Mb or 100Mb), along with integrated management software and bundle support. This architecture allows the building of an initial configuration with HPE compute and storage servers, including heterogeneous compute and storage in the same cluster, purpose-matched for your planned workload with the ability to add the correct node profile as needs change. The architecture was initially focused on building blocks for the core data center but is now expanding to include the edge and AI as well.

Benefits of an elastic design:

- Independently scale compute and storage nodes as business requirements change
- Run multiple workloads on the same cluster
- Reduce cost and security ramifications of having multiple copies of the data on isolated workload clusters
- Flexibility to repurpose compute nodes to support new services and application requirements

Hewlett Packard Enterprise also supports traditional Hadoop deployments that scale compute and storage together (symmetric), with some flexibility in choice of memory, processor, and storage capacity. This is widely based on the HPE ProLiant DL380 server platform, with density optimized architectures utilizing HPE Apollo 4200 series servers. Additionally, organizations that have already invested in balanced symmetric configurations have the ability to repurpose existing deployments into an HPE elastic architecture. This can help address scenarios where customers are looking to expand analytics capabilities by growing compute and/or storage capacity independently, without building a new cluster.

Figure 2 highlights some of the different server building blocks within the HPE Elastic Platform for Analytics, categorized by edge and core architectures and highlighting consumption-based service offerings from HPE. By leveraging a building block approach, customers can simplify the underlying infrastructure needed to address a myriad of different business initiatives around Data Warehouse modernization, analytics and BI, and build large-scale data lakes incorporating diverse sets of data. As workloads/compute and data storage requirements change (each often growing uncorrelated to the other) the HPE EPA elastic architecture framework allows customers to easily scale by adding compute and storage blocks independently.

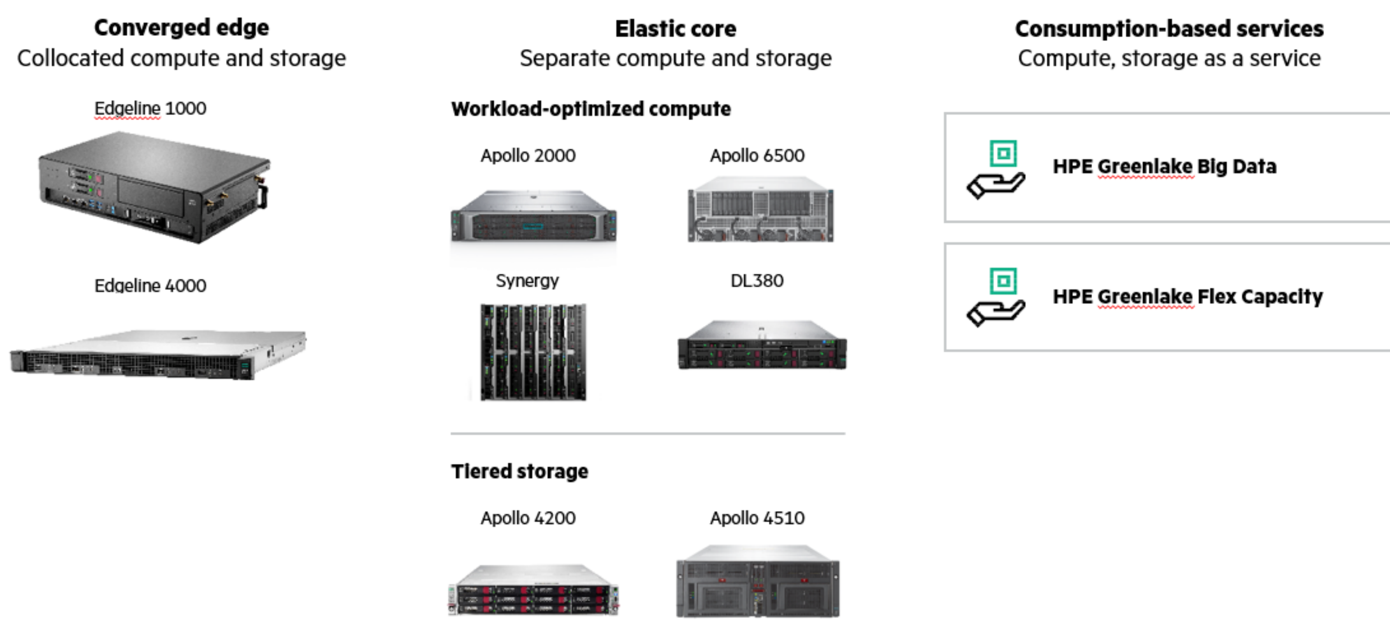

Figure 2. HPE products and services for the data pipeline

Solution overview

At the core of the HPE Elastic Platform for Analytics framework are the underlying workload building blocks. These blocks form the foundation of the EPA solution designs and can be pieced together in different ways to solve unique customer needs. This section will detail storage and compute blocks within elastic and traditional architectures. At the end of each section additional detail regarding multi-rack scalability of each design will be provided. The blocks defined in this section may be modified to address new or changing workloads and environments.

HPE Elastic Platform for Analytics architecture design

As Hadoop adoption expands in the enterprise, it is common to see multiple clusters running different workloads on a variety of technologies and Hadoop distributions, in both development and production environments, leading to challenges with data duplication and cluster sprawl. The HPE EPA architecture allows for consolidation of data and isolated workload clusters on a shared and centrally managed platform, providing organizations the flexibility to scale compute and/or storage as required, using modular building blocks for maximum performance and density. This fundamentally changes the economics of the solution across scale, performance and cost efficiency to meet specific use case and workload needs. HPE EPA building blocks provide the flexibility to optimize each workload and access the same pool of data for batch, interactive and real-time analytics.

By utilizing this elastic architecture, organizations can build and scale their infrastructure in alignment with the different phases and requirements of the existing analytics implementation. For example, some pilot environments might deploy a smaller number of symmetric configured storage-density optimized servers for data staging and basic MapReduce workloads; high-latency compute nodes can then be added (repurposing existing symmetric nodes to storage nodes through YARN labels), effectively migrating to an asymmetric architecture, without adding additional storage. A low-latency tier of density-optimized analytics blocks with SSDs or NVMe flash can be added to address the requirements of time series analysis of large datasets in real time. By adding workload-specific nodes and adjusting the compute-to-storage-node ratio by the block, organizations have the ability to tune their cluster towards business requirements, at the rack level. This architecture truly maximizes modern infrastructure density, agility, efficiency and performance while minimizing time-to-value, operational expenses, valuable data center real estate, and purchases of unneeded resources in order to meet mission requirements.

With storage and compute capacity growing at different rates for modern workloads, the HPE EPA architecture provides density optimized building blocks to target a variety of latency, capacity, and performance requirements.

The HPE EPA solution infrastructure blueprints are composed of five blocks: compute blocks, storage blocks, control blocks, network blocks and rack blocks. In addition to the five standard blocks that typically comprise an elastic solution, HPE has also developed blocks for optimizing workload performance, storage efficiency and deployment. Different than compute and storage blocks, optional dense compute or dense storage blocks can be combined to address various requirements (hot/cold storage, high/low latency compute, NoSQL, deep learning, etc.). Listed below are additional features of the blocks available in an HPE EPA solution:

- The standard compute block is one HPE Apollo 2000 chassis consisting of four HPE ProLiant XL170r Gen10 servers.
- The standard storage block is one HPE Apollo 4200 Gen10, consisting of 28x LFF HDDs or SSDs.
- The control block is made up of three HPE ProLiant DL360 Gen10 servers, with an optional fourth server acting as an edge or gateway node depending on the customer enterprise network requirements.
- The elastic network block consists of two HPE FlexFabric 5950 48SFP28 8QSFP28 switches and one HPE FlexFabric 5900AF-48G-4XG-2QSFP+ 1Gb switch.
- The aggregation network block consists of two HPE FlexFabric 5950 32QSFP28 (1U) switches. Used when adding a third rack.
- The rack block consists of either a 1200mm or 1075mm rack and its accessories.
- Examples of workload-optimized blocks include:
  - HPE Edgeline EL4000 Converged Edge system with HPE ProLiant m510 server cartridges for compute
  - Apollo 2000 with ProLiant XL170r Gen10 server with 512GB of memory
  - Apollo 6500 Gen10 server with GPUs
  - Apollo 4200 with 6TB, 8TB, 10TB or 12TB LFF HDDs
  - Apollo 4510 with 4TB, 6TB, 8TB, 10TB or 12TB LFF HDDs

Figure 3 provides a basic conceptual diagram of the HPE elastic architecture utilizing different block types.

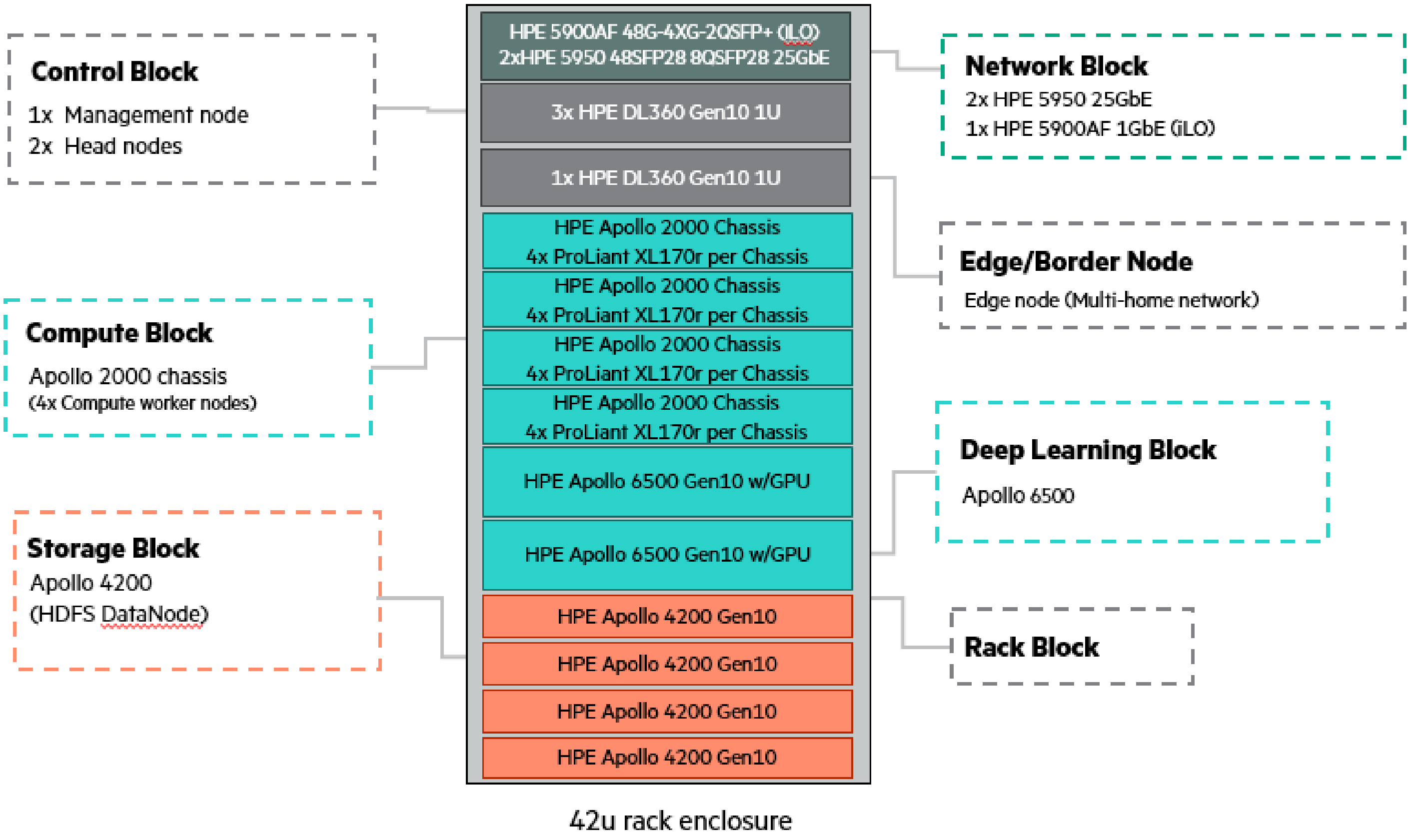

Figure 3. HPE EPA conceptual diagram

HPE EPA standard blocks – Apollo based

HPE Apollo 2000 compute block

The Apollo 2000 provides a density optimized compute platform that doubles the density of traditional 1U servers, providing up to four 2P servers in a 2U form factor. The HPE ProLiant XL170r Gen10 server is an ideal node for compute-intensive Spark, Hive/Tez, MapReduce, YARN, Vertica® SQL on Hadoop, and/or in-memory analytics in real-time, interactive, and/or batch workloads.

Table 1 provides the standard Apollo 2000 Gen10 compute block for the HPE elastic architecture.

Table 1. HPE Elastic Platform for Analytics – elastic compute block – Apollo 2000 Gen10

Component	Recommended configuration
Model	HPE Apollo r2600 Gen10 24SFF chassis with 4x HPE XL170r Gen10 servers
Processor	(2) Intel® Xeon®-Gold 6130 (2.1GHz/16-core/125W) processors
Memory	384GB RAM
OS disks	(2) HPE 480GB SATA 6Gb MU SC DS SSD or equivalent
Controller	HPE Smart Array E208i-p SR Gen10 12G SAS PCIe Plug-in Controller
Data disks	(2) HPE 480GB SATA 6Gb MU SC DS SSD
Network card	HPE Eth 10/25Gb 2P 640FLR-SFP28 Adapter

HPE Apollo 4200 storage block

HPE Apollo 4200 servers are cost-effective industry-standard storage servers, developed for Big Data with converged infrastructure that offers high density and energy-efficient storage. They perform as elastic storage nodes for HDFS and allow organizations to build multi-temperature data lake environments. The HPE Apollo 4200 is the foundation storage block for workload consolidation and is a key platform in the migration from a traditional to an elastic architecture.

Table 2 provides the standard Apollo 4200 storage block for the HPE EPA architecture.

Table 2. HPE Elastic Platform for Analytics – elastic storage block – Apollo 4200

Component	Recommended configuration
Model	HPE Apollo 4200 Gen10 24LFF server, plus optional LFF rear 4HDD cage
Processor	(2) Intel Xeon-Gold 5118 (2.3GHz/12-core/105W) processors
Memory	192GB RAM
OS disks	(2) HPE 480GB SATA MU M.2 2280 DS SSD or equivalent
Controller	(2) HPE Smart Array E208i-a/p SR G10 Ctrlr
Data disks	(28) HPE 4TB 6G SATA 7.2K LFF MDL LP HDD
Network card	HPE Eth 10/25Gb 2P 640SFP28 Adptr

The HPE Apollo 4200 allows you to save valuable data center space through its unique density-optimized 2U form factor that holds up to 28 LFF disks; and, with capacity for up to 336TB per node, this node will scale with Big Data environments.

HPE EPA standard blocks – HPE ProLiant DL based

HPE ProLiant DL360 Gen10 compute block

In addition to being the standard HPE EPA control block server, the HPE ProLiant DL360 Gen10 can also act as a compute platform that doubles the density of traditional 2U Big Data worker node servers. The HPE ProLiant DL360 Gen10 server is an ideal node for compute-intensive Spark, Hive/Tez, MapReduce, YARN, Vertica SQL on Hadoop, and/or in-memory analytics in real-time, interactive, and/or batch workloads.

Additionally, organizations that have already invested in balanced symmetric configurations (typically based on HPE ProLiant DL380 worker node platforms) have the ability to repurpose existing deployments into an HPE EPA architecture. Leveraging these symmetric clusters as a storage tier, and adding compute-dense platforms such as the HPE ProLiant DL360 Gen10 can help address scenarios where customers are looking to expand analytics capabilities by growing compute and/or storage capacity independently, without building a new cluster.

For more information see the, HPE Reference Configuration for the Elastic Platform for Big Data Analytics – ProLiant DL300-series modular building blocks of compute and storage optimized for modern workloads document at, http://h20195.www2.hpe.com/V2/GetDocument.aspx?docname=a00005869enw.

Table 3 provides the ProLiant DL360 Gen10 compute block for the HPE elastic architecture.

Table 3. HPE Elastic Platform for Analytics – elastic compute block – HPE ProLiant DL360 Gen10

Component	Recommended configuration
Model	HPE ProLiant DL360 Gen10 8SFF CTO server
Processor	(2) Intel Xeon-Gold 6130 (2.1GHz/16-core/125W) processors
Memory	384GB RAM
OS disks	(2) HPE 480GB SATA 6Gb MU SC DS SSD
Controller	HPE Smart Array P408i-a SR Gen10 12G SAS PCIe Plug-in Controller
Data disks	--
Network card	HPE Eth 10/25Gb 2P 640FLR-SFP28 Adapter

HPE ProLiant DL380 Gen10 storage block

HPE ProLiant DL380 Gen10 servers are cost-effective industry-standard servers and can serve as an alternative worker node. The HPE ProLiant DL380 performs as elastic storage nodes for HDFS and allows organizations to build multi-temperature data lake environments. The HPE ProLiant DL380 Gen10 can be leveraged as a storage block for low latency workloads, and for environments where failure domain considerations prevent higher density storage nodes.

Table 4 provides the HPE ProLiant DL380 Gen10 compute block for the HPE EPA architecture.

Table 4. HPE Elastic Platform for Analytics – EPA compute block – HPE ProLiant DL380 Gen10

Component	Recommended configuration
Model	HPE ProLiant DL380 Gen10 12LFF server, plus optional LFF rear 3HDD cage and LFF 4HDD Midplane cage
Processor	(2) Intel Xeon-Gold 5118 (2.3GHz/12-core/105W) processors
Memory	192GB RAM
OS disks	(2) 150GB SATA RI M.2 DS SSD or equivalent
Controller	HPE Smart Array P816i-a SR Gen10 controller, plus SAS Expander or Smart Array E208i-p
Data disks	(19) HPE 4TB 6G SATA 7.2k 3.5in SC MDL HDD
Network card	HPE Eth 10/25Gb 2P 640FLR-SFP28 Adapter

HPE EPA architecture workload-optimized blocks

HPE Converged IoT compute block – Edgeline EL4000 – with m510 server cartridge

HPE Edgeline EL4000 is one of the industry's first "Converged Edge Systems", integrating unprecedented levels of edge compute, precision data capture and control, data center class security, device and systems managements, as well as large and blazingly fast storage capability - all in one converged box. This rugged and compact system is designed to perform in harsher edge environments with higher shock, vibration and temperature levels. It is perfect for expanding your IoT infrastructure beyond traditional data center confines and can also help break cloud vendor lock-in to enable true edge computing. HPE Edgeline Converged Edge Systems are also designed to greatly accelerate the speed at which IoT data can be acquired, analyzed and actioned upon. This enables your business to make real-time decisions that can add value to your operations and result in better business outcomes.

Table 5 provides the converged IoT m510 compute block for the HPE EPA architecture.

Table 5. HPE Elastic Platform for Analytics – EPA accelerated compute block – low latency NoSQL – m510

Component	Recommended configuration
Model	HPE Edgeline EL4000 chassis with (4) HPE ProLiant m510 server cartridges
Processor	Intel Xeon D-1587 (1.7GHz/16-core/65W) processor
Memory	128GB RAM
OS disks	HPE 240GB M.2 SATA 2242 SSD or equivalent
Controller	SATA and PCIe controllers are Integrated in the Intel Xeon D SoC
Data disks	(2) 1024GB M.2 NVMe 2280 SSD
Network card	Mellanox Connect-X3 Pro Dual 10GbE

HPE Apollo 6500 Gen10 – Deep learning compute block with GPU cartridges

Deep learning is a subset of a more general field of artificial intelligence known as machine learning, using algorithms that extract some sort of structured data from huge amounts of unsupervised data. These algorithms often require lots of compute power (CPU and GPU). The HPE Apollo 6500 Gen10 System is an ideal HPC and Deep Learning platform providing unprecedented performance with industry leading GPUs, fast GPU interconnect, high bandwidth fabric and a configurable GPU topology to match your workloads. Ideal for deep learning workloads, the HPE Apollo 6500 Gen10 System is suitable for high performance computing workloads such as simulation and modelling. Eight GPU per server offers a faster and more economical deep learning system compared to more servers with fewer GPU each, keeping researchers productive as they iterate on a model more rapidly for a better solution in less time. HPC and AI models that would consume days or weeks can now be trained in a few hours or minutes.

Table 6 provides the Apollo 6500 Gen10 compute block with NVIDIA® GPUs for the HPE elastic architecture.

Table 6. HPE Elastic Platform for Big Data Analytics – elastic compute block – deep learning – Apollo 6500 Gen10

Component	Recommended configuration
Model	HPE Apollo 6500 Gen10 system
Processor	(2) Intel Xeon-Gold 6130 (2.1GHz/16-core/125W) processors
Memory	768GB RAM
OS disks	(2) HPE 480GB SATA 6Gb MU SC DS SSD or equivalent
Controller	HPE Smart Array P408i-a SR Gen10 Controller
Data disks	(2) HPE 480GB SATA 6Gb MU SC DS SSD
Graphics option	(8) HPE NVIDIA Tesla P100 SXM2 16GB Module
Network card	HPE Eth 10/25Gb 2P 640FLR-SFP28 Adapter

HPE Apollo 2000 – In-Memory analytics block

For workloads that require additional memory capacity, the in-memory analytics accelerator block provides additional memory capacity per node, 768GB versus 384GB in the standard compute block.

Table 7 provides the accelerated Apollo 2000 Gen10 compute block for the HPE elastic architecture.

Table 7. HPE Elastic Platform for Big Data Analytics – elastic accelerated compute block – in-memory analytics – Apollo 2000 Gen10

Component	Recommended configuration
Model	HPE Apollo r2600 Gen10 24SFF chassis with 4x HPE XL170r Gen10 servers
Processor	(2) Intel Xeon-Gold 6130 (2.1GHz/16-core/125W) processors
Memory	768GB RAM
OS disks	(2) HPE 480GB SATA 6Gb MU SC DS SSD or equivalent
Controller	HPE Smart Array E208i-p SR Gen10 12G SAS PCIe Plug-in Controller
Data disks	(2) HPE 480GB SATA 6Gb MU SC DS SSD
Network card	HPE Eth 10/25Gb 2P 640FLR-SFP28 Adapter

HPE Apollo 4510 dense storage block

The standard storage blocks outlined above are modeled to provide a balance between disk capacity, performance, and recovery times. For configurations where storage density is a driving factor, there is the HPE Apollo 4510. The HPE Apollo 4510 can provide an ideal platform for cold or archived storage and for customers that are looking to implement hot/cold data temperature zones into the cluster design. It is imperative that the networking backbone is sufficient to provide necessary data transfer rates and bandwidth. When choosing disk capacity it is important to analyze not only the cost of the drives but also the impact of recovery time in case of a failure.

The HPE Apollo 4510 provides an ultra-dense storage platform for archival storage, supporting up to 60 Large Form Factor (LFF) drives in a 4U chassis for a maximum of 720TB per system. The HPE Apollo 4510 servers include two sockets using the Intel Xeon-Gold series processors, DDR4-2666MHz memory and multiple HPE Smart Array options for internal bulk storage devices.

Table 8 provides the Apollo 4510 Gen10 storage block for the HPE EPA architecture.

Table 8. HPE Elastic Platform for Analytics – elastic dense storage block – Apollo 4510 Gen10

Component	Recommended configuration
Model	HPE Apollo 4510 Gen10 60LFF server
Processor	(2) Intel Xeon-Gold 5115 (2.4GHz/10-core/85W) processors
Memory	192GB RAM
OS disks	(2) HPE 1TB SATA 6G Midline 7.2k 2.5in SC HDD
Controllers	(2) HPE Smart Array P408i-p SR Gen10 Controller, HPE Smart Array E208i-a SR Gen10 Controller
Data disks	(60) HPE 4TB SATA 6G Midline 7.2k 3.5in LP HDD
Network card	(2) HPE Eth 10/25Gb 2P 640FLR-SFP28 Adapter

HPE Apollo 2000 or HPE DL360 compute blocks with BlueData EPIC Software

The HPE Elastic Platform for Big Data Analytics (EPA) harnesses the power of fast Ethernet networks and density optimized servers, providing a scalable and elastic multi-tenant architecture for Big-Data-as-a-Service when combined with the BlueData EPIC software platform. The BlueData EPIC (Elastic Private Instant Clusters) software platform makes it easier, faster, and more cost-effective for enterprises to deploy big data infrastructure and applications. When combined with the modern Big Data architecture provided by HPE Elastic Platform for Big Data Analytics, BlueData EPIC provides the flexibility and agility to support the rapidly changing requirements and uses cases for Big Data analytics, providing the ability to spin up and down Hadoop services as needed, creating a cloud-like experience and maximizing the use of on-premises infrastructure.

With BlueData EPIC and the HPE Elastic Platform for Big Data Analytics, enterprises can now deliver value from Big Data within days instead of months and with significantly lower TCO compared to traditional approaches. They can deliver the self-service agility, elasticity, and flexibility of Big-Data-as-a-Service in an on-premises deployment model.

For more information regarding the HPE Elastic Platform for Big Data Analytics and BlueData EPIC software, refer to the HPE Reference Configuration for Big-Data-as-a-Service with BlueData EPIC software document at, http://h20195.www2.hpe.com/V2/GetDocument.aspx?docname=4AA6-7823ENW.

HPE Synergy analytics blocks with Mesosphere DC/OS

With the advent of containers to virtualize workloads at scale without the associated performance penalty of virtual machines compared to bare metal, enterprises are exploring ways to design modern micro service based applications. However, one of the many challenges for the enterprise, in a rapidly evolving container space, is the fact that legacy and big data applications are stateful, and persistent; whereas, containerized applications typically are stateless and transient in nature. As many organizations attempt to build their own containerized big data environments with open source container orchestration tools like Kubernetes and Mesos with Mesosphere DC/OS, they soon realize that there are many other components like security, networking, build management, persistent storage, etc. that need to be integrated manually to have a functioning big data environment. At the same time, the evolution of big data based on Hadoop to fast data based on Spark enables separation of the data processing and data storage layers without being tied into a specific resource management framework like YARN, which allows for greater freedom not only in selecting and scaling of compute and storage, but also in the selection of resource management frameworks like Apache Mesos.

When faced with these challenges, enterprises tend to typically migrate towards a public cloud model that allows the infrastructure to be provisioned on-demand with all the necessary components required for building a data processing pipeline. However, data gravity tends to dictate a hybrid approach to big data, with some workloads in the cloud and others on-premises or at the edge.

To address these needs, Hewlett Packard Enterprise and Mesosphere have developed a joint solution leveraging Mesosphere Enterprise DC/OS software and the HPE Elastic Platform for Big Data Analytics to build a cost-effective and flexible on-premises architecture for deploying scalable Fast Data analytics workloads.

Mesosphere Enterprise DC/OS is an enterprise grade, data center operating system, providing a platform for securely provisioning containers at production scale to support next generation analytics workloads and frameworks.

Mesosphere DC/OS is based on the production proven Apache Mesos, and allows operators and data scientists to accelerate deployment and operation of advanced data services while maximizing utilization and reducing infrastructure and cloud cost. The HPE Elastic Platform for Big Data Analytics is a modular and logical framework that allows customers to choose the right blocks for the workload to achieve optimum performance and speed of deployment, while reducing big data cluster footprint by optimizing for density.

The combination of HPE EPA architecture with Analytics blocks built on HPE Synergy and Mesosphere DC/OS (Data Center / Operating System) is designed to deliver all the benefits of the cloud, enabling customers to rapidly provision Fast Data infrastructure on-premises with the benefits of a secure, scalable, and high performance architecture for Fast Data Analytics. This can dramatically reduce deployment complexity while improving business agility by providing an elastic self-service infrastructure for the big data services made available via the Mesosphere Catalog. The time-to-value for big data deployments can be reduced from weeks to days, while reducing overall costs compared to traditional bare-metal deployments.

For an in-depth analysis of the HPE Synergy analytics blocks for the EPA architecture with Mesosphere DC/OS solution, please see the HPE Reference Architecture for Fast Data Analytics on Mesosphere DC/OS technical white paper at https://h20195.www2.hpe.com/v2/getdocument.aspx?docname=a00039256enw.

Multi-rack configuration

The single-rack HPE EPA architecture is designed to perform well as a standalone solution, and also form the basis of a multi-rack solution, as shown in Figure 4. When moving from a single-rack to a multi-rack solution, you simply add racks without having to change any components within the base rack. For performance and redundancy, two HPE FlexFabric 5950 48SFP28 8QSFP28 ToR switches are specified per expansion rack. The HPE FlexFabric 5950 48SFP28 8QSFP28 switch includes up to eight 100GbE uplinks that can be used to connect these switches to the desired network via a pair of HPE FlexFabric 5950 48SFP28 8QSFP28 aggregation switches. Depending on the quantity of ports required for customer connectivity and the bandwidth needed between racks it is possible for a single pair of aggregation switches to support up to 16 racks.

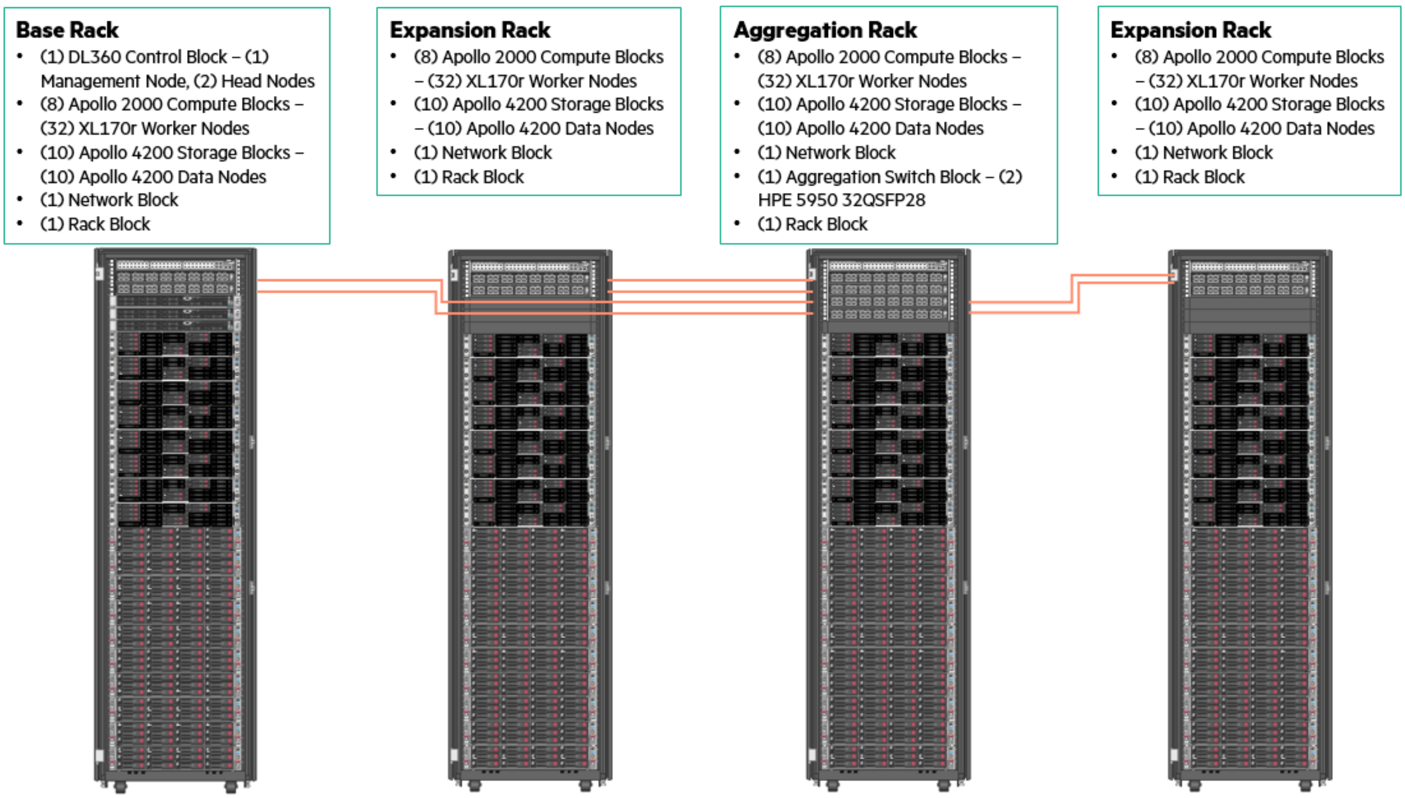

Figure 4. Multi-rack HPE EPA architecture, extending the capabilities of a single rack

A multi-rack solution assumes the base rack is already in place and extends its scalability. For example, the base rack already provides sufficient management services for a large scale-out system. Much of the architecture for the multi-rack solution is based on previous iterations of Hadoop testing, and is provided here as a general guideline for designing multi-rack Hadoop clusters.

HPE traditional architecture design

For conventional cluster models, an HPE traditional architecture leveraging the HPE ProLiant DL380 Gen10 server can provide balanced compute/storage, while an Apollo 4200 can provide exceptional storage density while being a key server for migration to an elastic architecture in the future. The HPE traditional architecture infrastructure blueprints are composed of four blocks: compute blocks, control blocks, network blocks and rack blocks. Listed below are additional features of the blocks available in a traditional architecture:

- The balanced compute block is one HPE ProLiant DL380 Gen10 server.
- The density optimized compute block is one Apollo 4200 Gen10, consisting of 28x LFF HDDs.
- The control block is made up of three HPE ProLiant DL360 Gen10 servers, with an optional fourth server acting as an edge or gateway node depending on the customer enterprise network requirements.

- The HPE traditional architecture network block consists of two HPE FlexFabric 5940-48XGT-6QSFP+.
- The aggregation network block consists of two HPE FlexFabric 5950-32QSFP28 (1U) switches. Used when adding a third rack.
- The HPE elastic architecture rack block consists of either a 1200mm or 1075mm rack and its accessories.

HPE traditional architecture standard blocks

HPE ProLiant DL380 balanced block

The HPE ProLiant DL380 Gen10 (2U) is the most widely deployed Hadoop worker node in conventional deployments, with some organization's utilizing clusters of the HPE ProLiant DL380 extending into the 1,000s of nodes. HPE ProLiant DL380 Gen10 servers deliver the best performance and expandability in the HPE 2P rack portfolio. The HPE ProLiant DL380 is known for its reliability, serviceability, near continuous availability and comprehensive warranty, making it the ideal node for most standard deployments. They have long been the server of choice for conventional Hadoop deployments worldwide as they provide the preferred CPU core-to-spindle ratio that has typically defined symmetric Hadoop worker node configurations.

The choice of embedded 4x1GbE, HPE FlexibleLOM or PCIe standup 1GbE to 100GbE adapters provide flexibility of networking bandwidth and fabric, allowing customers to adapt and grow to support changing business needs.

Table 9 provides the HPE ProLiant DL380 Gen10 compute block for the HPE traditional architecture.

Table 9. HPE traditional cluster worker node block

Component	Recommended configuration
Model	HPE ProLiant DL380 Gen10 12LFF server, plus optional LFF rear 3HDD cage and LFF 4HDD Midplane cage
Processor	(2) Intel Xeon-Gold 6130 (2.1GHz/16-core/125W) processors
Memory	384GB RAM
OS disks	(2) 150GB SATA RI M.2 DS SSD or equivalent
Controller	HPE Smart Array P816i-a SR Gen10 controller, plus SAS Expander or Smart Array E208i-p
Data disks	(19) HPE 4TB 6G SATA 7.2k 3.5in SC MDL HDD
Network card	HPE 10Gb 2-Port 562FLR-T

HPE Apollo 4200 – Density Optimized block

HPE Apollo 4200 servers are cost-effective industry–standard 2U storage-optimized servers, purpose built for Big Data with converged infrastructure and ProLiant technology that offers high density energy-efficient storage at up to 280TB per node, or 5.6PB per rack. Ideal for EDW offload workloads, the HPE Apollo 4200 is the foundation storage block for workload consolidation and is a key platform in the migration from a traditional to an elastic architecture.

Table 10 provides the standard Apollo 4200 density-optimized compute block for the HPE traditional architecture.

Table 10. HPE traditional cluster density-optimized worker node block – Apollo 4200

Component	Recommended configuration
Model	HPE Apollo 4200 Gen10 24LFF server, plus optional LFF rear 4HDD cage
Processor	(2) Intel Xeon-Gold 6132 (2.6GHz/14-core/140W) processors
Memory	384GB RAM
OS disks	(2) HPE 480GB SATA MU M.2 2280 DS SSD or equivalent
Controller	(2) HPE Smart Array E208i-a/p SR G10 Ctrlr
Data disks	(28) HPE 4TB 6G SATA 7.2K LFF MDL LP HDD
Network card	HPE 10Gb 2-Port 535T

Multi-rack configuration

The single-rack HPE traditional architecture is designed to perform well as a standalone solution, and also form the basis of a much larger multi-rack solution, as shown in Figure 5. When moving from a single-rack to a multi-rack solution, you simply add racks without having to change any components within the base rack.

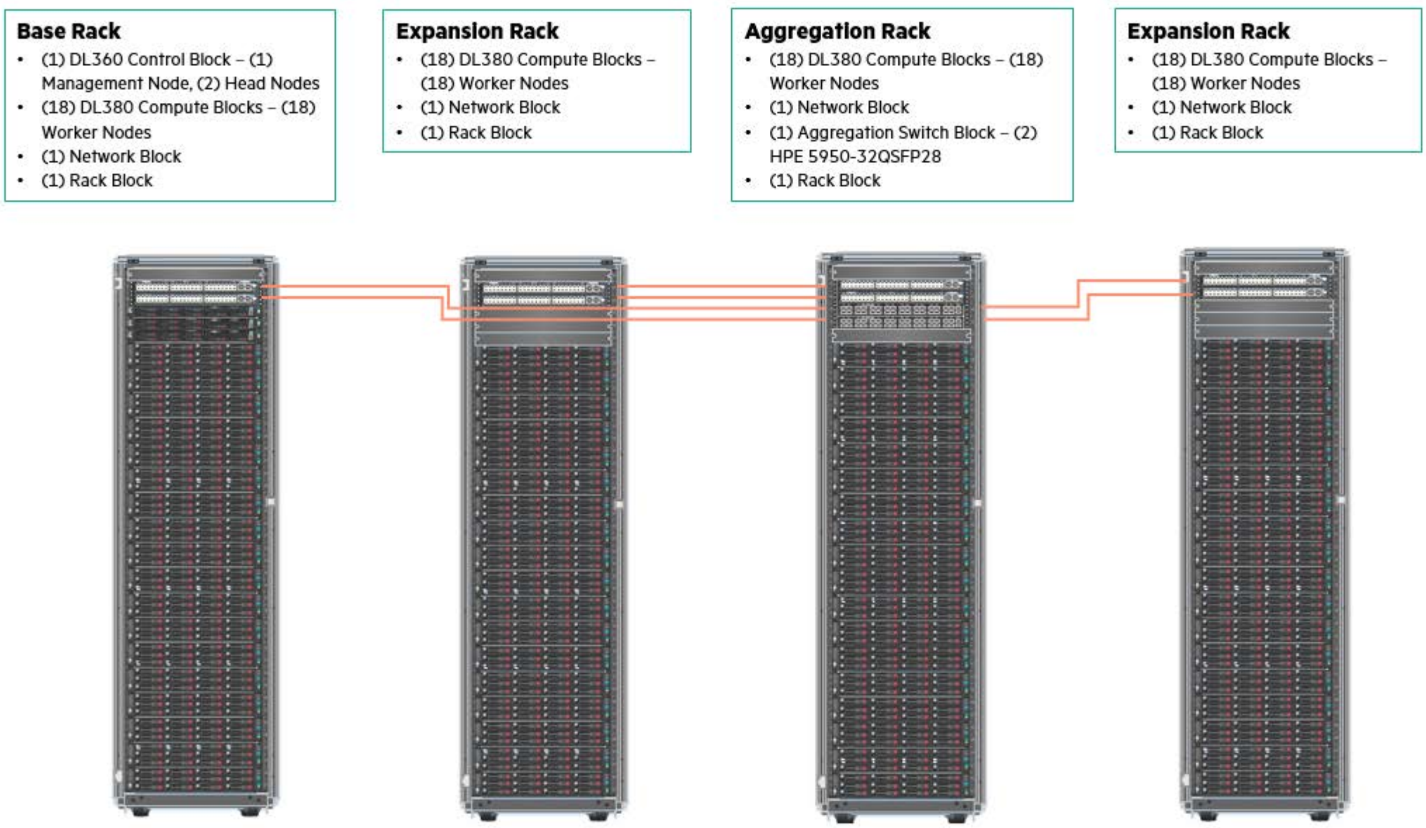

Figure 5. Multi-rack reference architecture (extension of the single-rack HPE traditional architecture)

The multi-rack design assumes the single-rack is already in place and extends its scalability. The single-rack configuration ensures the required amount of management services are in place for large scale out. For multi-rack clusters, one simply adds extension racks of a similar configuration to the single-rack configuration as shown in Figure 5. The expansion rack contains eighteen HPE ProLiant DL380 Gen10 servers and two HPE FlexFabric 5940-48XGT-6QSFP+ switches within a 42U rack; an additional 2x HPE ProLiant DL380 servers (2U each) or, two HPE FlexFabric 5950-32QSFP28 (1U) aggregation switches can be installed in the first expansion rack.

Note

There is no need for an aggregation switch when a second rack is added to the base rack. The existing HPE FlexFabric 5940 switches in the racks can be used to network between both (base racks) as shown in Figure 5. Figure 5 shows the network connection, using an additional aggregation switch, when a third rack (expansion rack) is added to the base racks.

HPE EPA control, network and rack blocks

HPE EPA control block

HPE ProLiant DL360 Gen10

HPE ProLiant DL360 Gen10 servers include two sockets using the Intel Xeon-Gold product family to provide the high performance required for management services on Hadoop clusters. The HPE EPA control block contains three HPE ProLiant DL360 Gen10 servers for use in both elastic and traditional architectures.

Table 11 provides the HPE ProLiant DL360 Gen10 control block for the HPE traditional and elastic architectures.

Table 11. HPE Elastic Platform for Analytics – control block – HPE ProLiant DL360 Gen10

Component	Recommended configuration
Model	HPE ProLiant DL360 Gen10 8SFF server
Processor	(2) Intel Xeon-Silver 4116 (2.1GHz/12-core/85W) processors
Memory	192GB RAM
OS disks	(2) HPE 1.2TB SAS 10K SFF SC DS HDD
Controller	HPE Smart Array P408i-a SR Gen10 controller
Data disks	(6) HPE 1.2TB SAS 10K SFF SC DS HDD
Network card	HPE Eth 10/25Gb 2P 640FLR-SFP28 Adapter (elastic) or HPE 10Gb 2-Port 562FLR-T (traditional)

A fourth HPE ProLiant DL360 Gen10 server may be configured with the control block to act as an edge node. An edge node acts as a gateway between the private VLAN of the cluster and the external routable network. Any application that requires both external network access and cluster private network access can run on this server. When significant storage bandwidth is required, use additional HPE ProLiant DL360 Gen10 servers as edge nodes.

HPE EPA network blocks

HPE EPA architecture network block

The HPE EPA architecture network block provides a high speed network that serves as the backbone of the environment. The switches were chosen after careful consideration and testing by HPE engineering and provide 25GbE networking to the storage nodes and compute nodes in the elastic architecture.

The HPE EPA architecture network block consists of two (2) IRF-bonded HPE FlexFabric 5950 48SFP28 8QSFP28 switches for high performance and redundancy. Each provides eight 100GbE uplinks that can be used to connect to the desired network, or, in a multi-rack configuration to connect to another pair of HPE FlexFabric 5950 48SFP28 8QSFP28 switches that are used for aggregation.

The HPE EPA architecture network block also includes a single HPE FlexFabric 5900AF-48G-4XG-2QSFP+ switch. Used exclusively to provide connectivity to HPE Integrated Lights-Out (iLO) management ports, which run at or below 1GbE. The iLO network is used for system provisioning and maintenance.

Table 12 provides the network block for the HPE elastic architecture.

Table 12. HPE Elastic Platform for Analytics – elastic network block

Component	Recommended configuration
Model	(2) HPE FlexFabric 5950 48SFP28 8QSFP28 switch
Model	HPE 5900AF-48G-4XG-2QSFP+ switch

Other network switches may be used in the HPE EPA architecture, as long as the bandwidth requirement is satisfied. For simplicity of this model the HPE FlexFabric 5950 48SFP28 8QSFP28 was chosen for the network block.

HPE traditional architecture network block

Different from the HPE EPA architecture network block, the HPE traditional architecture network block consists of two (2) IRF-bonded HPE 5940-48XGT-6QSFP+ 10GbE switches for high performance and redundancy. Chosen for their ability to provide sufficient port count and throughput, the HPE FlexFabric 5900AF-48XGT switches also include six 100GbE uplinks which can be used to connect the switches in the rack into the desired network. IRF bonding requires 2x 100GbE ports per switch, leaving 4x 100GbE ports on each HPE FlexFabric 5940-48XGT-6QSFP+ switch for uplinks.

A dedicated management switch for iLO traffic is not required as the HPE ProLiant DL360 Gen10, HPE ProLiant DL380 Gen10, and HPE Apollo 4200 Gen9 are able to share iLO traffic over NIC 1. The volume of iLO traffic is minimal and does not degrade performance over that port. Customers who would like to separate iLO and PXE traffic from the data/Hadoop network traffic can add a 1GbE HPE FlexFabric 5900AF-48G-4XG-2QSFP+ network switch.

Table 13 provides the network block for the HPE traditional architecture.

Table 13. HPE Elastic Platform for Analytics – traditional network block

Component	Recommended configuration
Model	(2) HPE 5940-48XGT-6QSFP+ switch

HPE EPA aggregation network block

In addition to the HPE elastic and traditional network blocks, HPE offers an aggregation network block consisting of two HPE FlexFabric 5950-32QSFP28 aggregation switches. The aggregation network block can be used with either elastic or traditional architectures. Depending on the quantity of ports required for customer connectivity and the bandwidth needed between racks it is possible for a single pair of aggregation switches to support up to 16 racks.

Table 14 provides the aggregation network block for the HPE EPA system.

Table 14. HPE Elastic Platform for Analytics – EPA aggregation network block

Component	Recommended configuration
Model	(2) HPE FlexFabric 5950 32QSFP28 switch

For more information regarding network architecture and design of both elastic and traditional architectures see, HPE Reference Configuration for networking best practices on the HPE Elastic Platform for Big Data Analytics (EPA) Hadoop ecosystem, at http://h20195.www2.hpe.com/V2/GetDocument.aspx?docname=a00004216enw

HPE EPA rack block

The HPE EPA rack block consists of a single rack and its accessories. Depending on the server types either a 1075mm or 1200mm deep rack may be used. Please consult your server documentation to determine if a deeper rack is required.

Table 15 provides the rack block for the HPE EPA system.

Table 15. HPE Elastic Platform for Big Data Analytics – EPA rack block

Component	Recommended configuration
Model	HPE 642 1200mm Shock Intelligent Rack or equivalent
Model	HPE 642 1075mm Shock Intelligent Series Rack or equivalent
PDU	PDU type and quantity depending on rack configuration and facility power

Configuration guidance for HPE Elastic Platform for Analytics (EPA)

The HPE Elastic Platform for Analytics provides the framework and infrastructure foundation needed to address the wide variety of challenges customers face when building Big Data solutions. By leveraging a building block model, elastic and traditional configuration blueprints can be designed for various customer workload needs. These blueprints can then be adapted to address unique workload or compute/capacity requirements that may arise. This section builds on the building block model introduced in the previous sections and provides several examples for traditional and EPA blueprints in half-rack and rack level increments.

The sizing for the configurations outlined in this section is driven by core storage requirements for the cluster. Storage capacity is often the core metric that is most reasonably understood at the start of a Big Data project, with compute performance requirements being less quantifiable until more in-depth testing of the specific customer workloads is understood. This lack of baseline compute performance information is one reason the flexibility of the EPA elastic blueprint mode is so important. As more knowledge of a customer's compute requirements are understood, the cluster compute tier can be scale-out without dragging along additional storage capacity that may not be warranted.

Hadoop cluster storage sizing requires careful planning and identifying the current and future storage and compute needs. Use the following as general guidelines for data inventory:

- Sources of data
- Frequency of data
- Raw storage
- Processed HDFS storage
- Replication factor
- Default compression turned on
- Space for intermediate files

To calculate your storage needs, determine the number of TB of data per day, week, month, and year; and then add the ingestion rates of all data sources.

- It makes sense to identify storage requirements for the short-, medium-, and long-term.
- Another important consideration is data retention – both size and duration. Which data must you keep? For how long?
- In addition, consider maximum fill-rate and file system format space requirements on hard drives when estimating storage size.

This capacity based modeling approach is one that is often used when sizing traditional clusters. Determining the quantity of worker nodes required in a traditional environment is a function of the total storage requirement divided by the storage capacity of the worker nodes (e.g., 19x 3TB HDD in an HPE ProLiant DL380 yields 57TB total raw capacity). Sizing of the storage nodes in an elastic environment is done in a similar fashion, but sizing of the compute nodes is based on the workload requirements and can be scaled accordingly.

The example configurations outlined in the remainder of this white paper cover the following four general scenarios:

- HPE Traditional Architecture – Standard
- HPE Traditional Architecture – Storage optimized
- HPE EPA Architecture – Standard
- HPE EPA Architecture – Storage optimized

These examples are intended to provide representable configuration blueprints aligned to the usable storage requirements for a Big Data solution. The key pivot points for using these designs is the choice between traditional and EPA models and the size of the cluster. The EPA examples can expand, depending on the requirements, to a variety of compute/storage nodes in a single rack, ranging from hot (with a large number of compute nodes and minimal storage) to cold (with a large number of storage nodes and minimal compute). Additional blocks can be added to the EPA framework to support NoSQL and HBase workloads, for example. The traditional examples can also expand, as long as CPU and memory demands scale in a constant manner.

The HPE EPA Sizing Tool, which models performance and sizing of various platforms in both traditional and elastic configurations based on testing against multiple workloads and platforms, serves as the source for the following comparison data. To download the HPE Sizer for the Elastic Platform for Analytics (HPE EPA Sizing Tool), visit hpe.com/info/sizers.

HPE traditional configurations – Standard

Table 16 provides example blueprints for standard traditional configurations using HPE ProLiant DL380 Gen10 servers, scaling from a POC starter kit up to a multi-rack configuration. Building block quantities are provided in parenthesis in the appropriate row. Usable capacity is determined using a replication level of 3, 1:1 compression ratio and overhead of 25%.

Table 16. Traditional configuration examples – HPE ProLiant DL380

	Half Rack	Full Rack	Multi Rack
Control block	(3) DL360	(3) DL360	(3) DL360
Compute block	(9) DL380 with 4TB HDD	(18) DL380 with 4TB HDD	(54) DL380 with 4TB HDD
Network block	(2) HPE 5940 switch	(2) HPE 5940 switch	(6) HPE 5940 switch
Aggregation network block			(2) HPE 5950 switch
Rack block	HPE 1075mm rack	HPE 1075mm rack	(3) HPE 1075mm rack
Specifications			
Raw disk capacity (TB)	684	1368	4104
Usable disk capacity (TB)	171	342	1026
Compute performance (SpecInt per U)	566.3	647.1	668.4
Storage performance (MB/s per U)	783.8	895.7	925.1

HPE traditional configurations – Storage optimized

Table 17 provides example blueprints for storage optimized traditional configurations. These configurations leverage the Apollo 4200 density optimized block in a traditional model. The specifications show approximately 85% more storage density per U when leveraging the Apollo 4200 compared to the HPE ProLiant DL380.

Usable capacity is determined using a replication level of 3, 1:1 compression ratio and overhead of 25%.

Table 17. Traditional configuration examples – Density Optimized Apollo 4200

	Half Rack	Full Rack	Multi Rack
Control block	(3) DL360	(3) DL360	(3) DL360
Compute block	(9) Apollo 4200 with 4TB HDD	(18) Apollo 4200 with 4TB HDD	(54) Apollo 4200 with 4TB HDD
Network block	(2) HPE 5900AF switch	(2) HPE 5900AF switch	(6) HPE 5900AF switch
Aggregation network block			(2) HPE 5950 switch
Rack block	HPE 1075mm rack	HPE 1075mm rack	(3) HPE 1075mm rack
Specifications			
Raw disk capacity (TB)	1008	2016	6048
Usable disk capacity (TB)	252	504	1512
Compute performance (SpecInt per U)	431.3	492.9	509
Storage performance (MB/s per U)	1155	1320	1363.3

HPE EPA configurations – Standard

Table 18 provides example blueprints for baseline EPA configurations using the standard elastic compute and storage blocks. Since the compute and storage tiers are disaggregated there is more compute capacity and storage density as compared to the standard traditional configurations. This is by design as the elastic model provides a foundation for running multiple disparate services without having to have separate standalone Hadoop clusters and the resulting duplication of data. Usable capacity is determined using a replication level of 3, 1:1 compression ratio and overhead of 25%.

Table 18. HPE EPA configuration examples – Standard elastic blocks

	Half Rack	Full Rack	Multi Rack
Control block	(3) DL360	(3) DL360	(3) DL360
Compute block	(16) XL170r	(32) XL170r	(96) XL170r
Storage block	(5) Apollo 4200 with 4TB HDD	(10) Apollo 4200 with 4TB HDD	(30) Apollo 4200 with 4TB HDD
Network block	(2) HPE 5950 switch (1) HPE 5900AF-48G switch	(2) HPE 5950 switch (1) HPE 5900AF-48G switch	(6) HPE 5950 switch (3) HPE 5900AF-48G switch
Aggregation network block			(2) HPE 5950 switch
Rack block	HPE 1075mm rack	HPE 1075mm rack	(3) HPE 1075mm rack
Specifications			
Raw disk capacity (TB)	575	1151	3452
Usable disk capacity (TB)	144	288	863
Compute performance (SpecInt per U)	1006.3	1027.7	1159.2
Storage performance (MB/s per U)	1355	1383.8	1561

HPE EPA configurations – Storage optimized

Table 19 provides example blueprints for half rack, full rack and multi-rack elastic configurations using the elastic density optimized Apollo 4200 storage block. This storage block leverages the 8TB LFF drives rather than the 4TB LFF which are part of the standard Apollo 4200 storage block. This provides the densest storage design but needs to be rationalized against customer failure domain and SLA requirements for node recovery. And while not shown in these baseline examples, the elastic model provides the ability to support hybrid storage models with different tiers of storage in the cluster. This allows operationally mature environments to configure different storage tiers and build data temperature constructs into the cluster design. Usable capacity is determined using a replication level of 3, 1:1 compression ratio and overhead of 25%.

Table 19. EPA configuration examples – Dense elastic Apollo 4200 block

	Half Rack	Full Rack	Multi Rack
Control block	(3) DL360	(3) DL360	(3) DL360
Compute block	(16) XL170r	(32) XL170r	(96) XL170r
Storage block	(5) Apollo 4200 with 8TB HDD	(10) Apollo 4200 with 8TB HDD	(30) Apollo 4200 with 8TB HDD
Network block	(2) HPE 5940 switch (1) HPE 5900AF-48G switch	(2) HPE 5940 switch (1) HPE 5900AF-48G switch	(6) HPE 5940 switch (3) HPE 5900AF-48G switch
Aggregation network block			(2) HPE 5950 switch
Rack block	HPE 1075mm rack	HPE 1075mm rack	(3) HPE 1075mm rack
Specifications			
Raw disk capacity (TB)	1135	2271	6812
Usable disk capacity (TB)	284	568	1703
Compute performance (SpecInt per U)	1006.3	1027.7	1159.2
Storage performance (MB/s per U)	1355	1383.8	1561

Summary

HPE Elastic Platform for Analytics offers an asymmetric architecture and a different approach to building a data pipeline optimized for Big Data Analytics. Consisting of separately definable and scalable compute and storage resources connected via high speed network components, this asymmetric architecture allows you to build an initial configuration with HPE compute and storage servers (including heterogeneous compute/storage in the same cluster) purpose matched for your workload with the ability to add the correct node profile as needs change.

Existing Big Data deployments using conventional symmetric architectures based on the HPE ProLiant DL380 can transition to a truly elastic platform with the HPE elastic architecture to reduce data center footprint, operating costs, optimize performance and efficiently manage the exploding growth of Big Data.

Appendix A: Gen9 Building Blocks

The following section contains information regarding HPE Gen9 traditional and elastic building blocks for reference purposes. These Gen9 servers will continue to be supported in the HPE Elastic Platform for Big Data Analytics.

HPE elastic architecture standard blocks

HPE Apollo 2000 Gen9 compute block

Table A-1 provides the standard Apollo 2000 Gen9 compute block for the HPE elastic architecture.

Table A-1. HPE Elastic Platform for Big Data Analytics – elastic compute block – Apollo 2000 Gen9

Component	Recommended configuration
Model	HPE Apollo r2600 24SFF chassis with 4x HPE XL170r Gen9 servers
Processor	(2) Intel Xeon E5-2680 v4 14-core processors
Memory	256GB RAM
OS disks	(2) HPE 480GB 6Gb SATA 2.5in MU-2 SC SSD or equivalent
Controller	HPE Dynamic Smart Array B140i controller
Data disks	(2) HPE 480GB 6Gb SATA 2.5in MU-2 SC SSD
Network card	HPE Eth 10/25Gb 2P 640FLR-SFP28 Adapter

HPE Apollo 4200 storage block

Table A-2 provides the standard Apollo 4200 storage block for the HPE elastic architecture.

Table A-2. HPE Elastic Platform for Big Data Analytics – elastic storage block – Apollo 4200

Component	Recommended configuration
Model	HPE Apollo 4200 Gen9 24LFF server, plus optional LFF rear 4HDD cage
Processor	(2) Intel E5-2660 v4 (2.0GHz/14-core/105W) processors
Memory	128GB RAM
OS disks	HPE Dual 120GB RI Solid State M.2 kit
Controller	HPE Smart Array P840ar/2G FIO controller
Data disks	(28) HPE 4TB 6G SATA 7.2k 3.5in MDL LP HDD
Network card	HPE Eth 10/25Gb 2P 640FLR-SFP28 Adptr

HPE Apollo 2000 – Deep learning accelerator block with XL190r Gen9 GPU cartridges

Table A-3 provides the accelerated XL190r compute block with dual NVIDIA GPUs for the HPE elastic architecture.

Table A-3. HPE Elastic Platform for Big Data Analytics – elastic accelerated compute block – deep learning – XL190r Gen9

Component	Recommended configuration
Model	HPE Apollo 2000 chassis with (2) HPE ProLiant XL190r server cartridges
Processor	(2) Intel E5-2680 v4 14-core processors
Memory	256GB RAM
OS disks	(2) HPE 480GB 6Gb SATA 2.5in MU-2 SC SSD or equivalent
Controller	HPE Dynamic Smart Array B140i controller
Data disks	(2) HPE 480GB 6Gb SATA 2.5in MU-2 SC SSD
Graphics option	HPE NVIDIA Tesla K80 Dual GPU module
Network card	HPE Eth 10/25Gb 2P 640FLR-SFP28 Adapter

HPE Apollo 2000 – In-Memory analytics accelerator block

Table A-4 provides the accelerated Apollo 2000 Gen9 compute block for the HPE elastic architecture.

Table A-4. HPE Elastic Platform for Big Data Analytics – elastic accelerated compute block – in-memory analytics – Apollo 2000 Gen9

Component	Recommended configuration
Model	HPE Apollo r2600 24SFF chassis with 4x HPE XL170r Gen9 servers
Processor	(2) Intel E5-2680 v4 14-core processors
Memory	512GB RAM
OS disks	(2) HPE 480GB 6Gb SATA 2.5in MU-2 SC SSD or equivalent
Controller	HPE Dynamic Smart Array B140i controller
Data disks	(2) HPE 480GB 6Gb SATA 2.5in MU-2 SC SSD
Network card	HPE Eth 10/25Gb 2P 640FLR-SFP28 Adapter

HPE traditional architecture standard blocks

HPE ProLiant DL380 Gen9 balanced block

Table A-5 provides the HPE ProLiant DL380 Gen9 compute block for the HPE traditional architecture.

Table A-5. HPE Elastic Platform for Big Data Analytics – traditional worker node block

Component	Recommended configuration
Model	HPE ProLiant DL380 Gen9 12LFF server, plus optional LFF rear 3HDD cage
Processor	(2) Intel E5-2680 v4 14-core processors
Memory	256GB RAM
OS disks	HPE Dual 340GB RI-2 Solid State M.2 kit
Controller	HPE Smart Array P840/4G controller
Data disks	(15) HPE 4TB 6G SATA 7.2k 3.5in SC MDL HDD
Network card	HPE 10Gb 2-Port 561FLR-T

HPE EPA control blocks

HPE ProLiant DL360 Gen9

Table A-6 provides the HPE ProLiant DL360 control block for the HPE elastic and traditional architectures.

Table A-6. HPE Elastic Platform for Big Data Analytics – control block – HPE ProLiant DL360

Component	Recommended configuration
Model	HPE ProLiant DL360 Gen9 8SFF server
Processor	(2) Intel E5-2640 v4 10-core processors
Memory	128GB RAM
OS disks	(2) HPE 900GB 12G SAS 10K 2.5in SC ENT HDD
Controller	HPE Smart Array P440ar/2G controller
Data disks	(6) HPE 900GB 12G SAS 10K 2.5in SC ENT HDD
Network card	HPE Eth 10/25Gb 2P 640FLR-SFP28 Adapter

HPE Apollo 2000 system with HPE ProLiant XL170r Gen9

Table A-7 provides the Apollo 2000 control block for the HPE elastic architecture.

Table A-7. HPE Elastic Platform for Big Data Analytics – WDO control block – Apollo 2000

Component	Recommended configuration
Model	HPE Apollo r2600 or r2800 24SFF chassis with 4x HPE XL170r servers
Processor	(2) Intel E5-2640 v4 10-core processors
Memory	128GB RAM
OS disks	(2) HPE 1.2TB 12G SAS 10K 2.5in SC ENT HDD
Controller	HPE Smart Array P440/4G controller
Data disks	(4) HPE 1.2TB 12G SAS 10K 2.5in SC ENT HDD (may vary if r2800 chassis is used)
Network card	HPE Eth 10/25Gb 2P 640FLR-SFP28 Adapter

Resources and additional links

HPE Reference Architectures, hpe.com/info/ra

HPE Apollo 2000 systems, hpe.com/us/en/servers/hpc-apollo-2000.html

HPE Apollo 4200 systems, hpe.com/us/en/servers/hpc-apollo-4000.html

HPE Apollo 6000 systems, hpe.com/us/en/servers/hpc-apollo-6000.html

HPE Moonshot system, hpe.com/us/en/servers/moonshot.html

HPE ProLiant DL380 servers, hpe.com/servers/DL380

HPE Networking, hpe.com/networking

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

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

4AA6-8931ENW, Rev. 4, February 2019