



Hewlett Packard
Enterprise

HPE Reference Architecture for SAP HANA Vora with Hortonworks HDP

Using HPE ConvergedSystem for SAP HANA and
HPE Elastic Platform for Big Data Analytics



HORTONWORKS®

Technical white paper

Contents

Executive summary.....	4
Introduction.....	4
Solution overview.....	5
Hardware overview.....	6
Software overview.....	8
SAP HANA Vora use cases.....	12
Data tiering.....	12
Analytics.....	12
Solution components.....	13
Hardware components.....	14
Software components.....	18
A proof-of-concept for interactive analytics and data tiering with SAP HANA Vora.....	20
Interactive analytics across SAP HANA and Hadoop.....	21
Data movement between SAP HANA and HDP / Apache Hadoop.....	24
Solution configuration guidance.....	24
Setting up SAP HANA Vora.....	24
SAP HANA Vora sizing guidelines.....	27
SAP HANA Vora query tuning guidelines.....	28
Capacity and sizing.....	28
Expanding the base configuration of HPE WDO system.....	28
Analysis and recommendations.....	30
Benefits of HPE Elastic Platform for Big Data Analytics (EPA) using SAP HANA Vora.....	30
Implementing a proof-of-concept.....	31
Summary.....	31
Appendix A: Bill of materials.....	31
Appendix B: Alternate HPE EPA components for compute and storage.....	36
Alternate compute node components.....	36
Alternate storage node components.....	37
Appendix C: HPE value added services and support.....	37
Factory Express Services.....	37
Technical Services Consulting – Reference Architecture Implementation Service for Hadoop (Hortonworks).....	37
Technical Services Consulting – Big Data Services.....	37
HPE Support options.....	38
Appendix D: Installing SAP HANA Vora on HPE EPA cluster.....	38
Installing HDP 2.3 and Spark.....	38
Installing and configuring SAP HANA Vora components.....	39

Appendix E: HPE appliances for SAP HANA details.....	57
HPE ConvergedSystem 500 for SAP HANA Scale-up and Scale-out appliances	57
Appendix F: SAP HANA configuration.....	58
Setting up SAP HANA to connect to HPE EPA cluster.....	58
Setting up SAP HANA Data Warehousing Foundation – Data Lifecycle Manager	58
Resources and additional links.....	60

Executive summary

As organizations struggle to manage and realize the value from exploding volumes of data, they seek more agile and capable analytic systems that speed up access to data that resides in multiple operational and Big Data stores.

Some of the drivers that target business value include improving customer retention, increasing operational efficiencies, influencing product development and quality, and gaining a competitive advantage. Other business drivers are more oriented towards cost reduction by simplifying data management, consolidation to multi-tenant platforms and intelligent data placement.

SAP HANA® Vora has emerged as a technology to bridge the divide between business critical enterprise data and exploding volumes of Big Data. SAP HANA Vora allows organizations to use familiar languages, such as SQL, to analyze and add insight across large volumes of operational and contextual data from enterprise applications, data warehouses, Hadoop data lakes and real time streams of data. In addition, SAP HANA Vora enables organizations to reduce costs of storing and processing Big Data by intelligently distributing data between Hadoop and SAP HANA based on speed and cost of data access, which minimizes large scale data movement and facilitates data analysis in place.

Hortonworks Data Platform (HDP) is the industry's only true secure, enterprise-ready open source Apache Hadoop distribution based on a centralized architecture (YARN). HDP addresses the complete needs of data-at-rest, powers real-time customer applications and delivers robust analytics that accelerate decision making and innovation.

This white paper describes the implementation of two key features of SAP HANA Vora with Spark:

- Interactive analytics across both SAP HANA and Hadoop data using SAP HANA Vora and Spark
- Manage data movement between SAP HANA and the Hortonworks Data Platform powered by Apache Hadoop

This implementation is based on Hortonworks Data Platform (HDP), 100% open source distribution of Apache Hadoop, specifically HDP 2.3 and the HPE Elastic Platform for Big Data Analytics (EPA). The configurations reflected in this document have been jointly designed and developed by HPE and Hortonworks to provide optimum computational performance for Hadoop.

Document purpose: The paper guides you through the deployment and configuration of SAP HANA Vora on the HPE Elastic Platform for Big Data Analytics using Hortonworks Data Platform solution. The paper also describes the HPE Apollo 2000 compute system and the HPE Apollo 4200 storage nodes as key components of the solution. In addition, this paper describes the best practices to integrate SAP HANA Vora and Hadoop with SAP HANA running on HPE ConvergedSystem for SAP HANA.

Target audience: The intended audience of this document includes, but is not limited to, IT managers, solution architects, sales engineers, services consultants, partner engineers and customers that are interested in deploying SAP HANA Vora for interactive analytics on Hadoop and SAP HANA on Hewlett Packard Enterprise and Hortonworks Data Platform (HDP) solutions.

This Reference Architecture describes the solution testing performed with SAP HANA Vora 1.2, Hortonworks HDP 2.3 and Apache Spark 1.5 in August 2016.

Introduction

According to industry surveys, almost 70% of available data in an enterprise lies un-analyzed due to a lack of the right tools and resources to quickly identify data of value from the ever increasing volumes and sources of raw data. Driving business outcomes in real time through deeper insight into data, by linking operational business data with Big Data from multiple sources, is a competitive necessity. This requires new approaches to combining and correlating operational data with raw data, from social feeds or sensor data, in a cost efficient and timely manner.

Apache Hadoop is a software framework that is being adopted by many enterprises as a cost-effective analytics platform for Big Data analytics. Hadoop is an ecosystem of several services rather than a single product, and is designed for storing and processing petabytes of data in a linear scale-out model. Each service in the Hadoop ecosystem may use different technologies to ingest, store, process and visualize data.

Hortonworks is a major contributor to Apache Hadoop, the world's most popular Big Data platform. Hortonworks focuses on further accelerating the development and adoption of Apache Hadoop by making the software more robust and easier to consume for enterprises and more open and extensible for solution providers.

Enterprises are now looking to deploy new real-time analytics use cases such as predictive analytics, fraud detection, recommendation engines, and advertising analytics. Spark has become a popular technology for these near real-time analytic workloads. In addition, Spark's data federation

capabilities enable the use of SQL to query across different data sources, allowing organizations to combine operational data stores, such as Oracle or SAP® enterprise data warehouses, with unstructured data in Hadoop.

While Spark and Hadoop offer the benefit of scaling out processing while cost-effectively managing petabytes of unstructured data, enabling unified data in-place access to both Hadoop data and data from transactional business systems, such as SAP Business Suite, or data warehouses such as SAP HANA, has been a challenge.

The second challenge relates to the cost of storing and processing exploding volumes of data. For SAP customers looking to lower the total cost of ownership of their SAP infrastructure, the SAP HANA platform aims to address this challenge, by consolidating the data management and processing of both transactional and analytics systems.

As SAP HANA evolves as a data platform to support new functionality and manage increasing volumes of data from new sources, lowering cost and simplifying data management have been key drivers for wider adoption. SAP HANA becomes expensive as data volumes move beyond a few terabytes, and may be overkill especially for storing raw or cold data.

Both legacy and modern applications are evolving towards architecture frameworks that are more open, and designed to run on standard infrastructure that can scale out in tandem with capacity demands. This has resulted in a more flexible and modular approach to infrastructure acquisition and cost-effective capacity expansion.

The solution described in this paper helps address the challenges of accelerating analytics on multiple sources of data, and helps lower the total cost of storing, managing and processing this data. SAP HANA Vora bridges the gap between enterprise data in SAP HANA and unstructured data in Hadoop. Using SAP HANA Vora in conjunction with Spark, Hadoop and SAP HANA, organizations can simplify their data access and management experience while accelerating their analytics workloads.

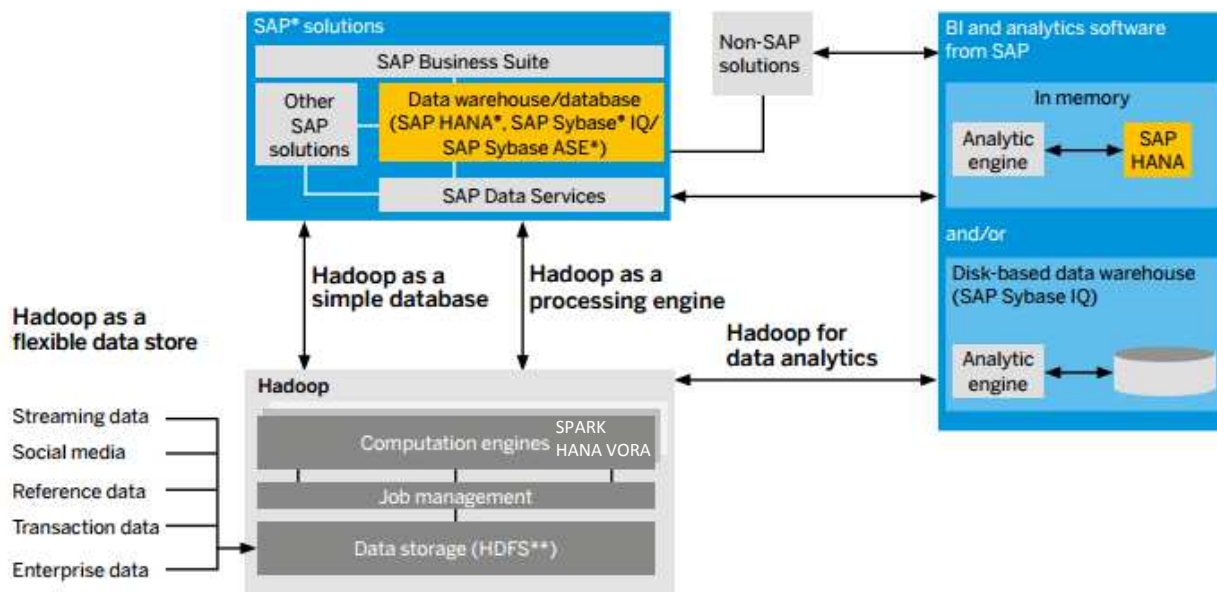


Figure 1. SAP HANA Vora overview (Source: SAP)

Solution overview

The HPE ConvergedSystem for SAP HANA and the HPE Elastic Platform for Big Data Analytics (EPA) provide you with a wide range of infrastructure building blocks to help you optimally deploy both SAP HANA and SAP HANA Vora to solve your business needs. These solutions simplify the infrastructure design and deployment phase for customers and provide an agile solution platform for helping customers drive innovation with new business models and outcomes.

The solution described in this document will focus on the deployment of SAP HANA Vora with Spark and Hadoop and its integration with SAP HANA for Analytics, and SAP Data Services using the Data Lifecycle Manager (DLM), on the following infrastructure solutions from HPE:

- HPE Elastic Platform for Big Data Analytics (EPA), running HDP with Spark, and SAP HANA Vora
- HPE ConvergedSystem 500 for SAP HANA (CS500), running SAP HANA and SAP DLM

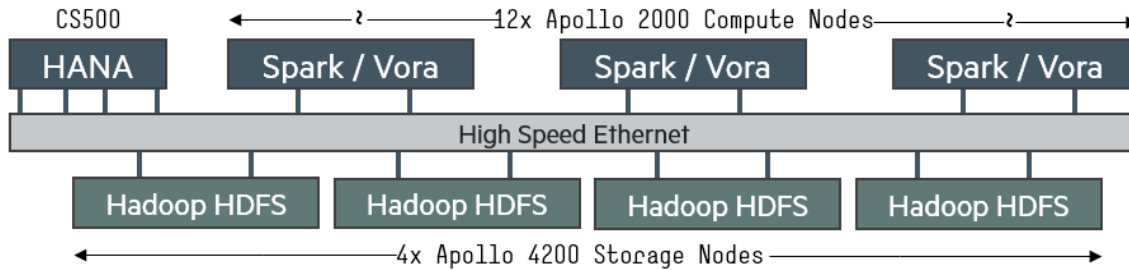


Figure 2. Solution overview - HPE ConvergedSystem 500 for SAP HANA and HPE Elastic Platform for Big Data Analytics (EPA), running HDP with SAP HANA Vora

Using a predictive analytics use case, we will detail the implementation of two key features of SAP HANA Vora with Spark:

- Interactive analytics across both SAP HANA and Hadoop data using SAP HANA Vora and Spark
- Manage data movement between SAP HANA and the Hortonworks Data Platform powered by Apache Hadoop

The implementation is based on Hortonworks’ Distribution of Apache Hadoop (HDP), specifically HDP 2.3 and Ambari 2.2, and the HPE Elastic Platform for Big Data Analytics system.

Hardware overview

HPE ConvergedSystem for SAP HANA

The HPE ConvergedSystem for SAP HANA is a portfolio of purpose built and optimized infrastructure solutions for SAP HANA. Available for use cases across side card analytics, data warehousing and SAP Business Suite on HANA / S/4HANA, the HPE ConvergedSystem for SAP HANA portfolio is an industry-leading SAP HANA infrastructure solution, offering customers mission critical performance, ease of manageability and support, and flexibility in consumption models.

The HPE ConvergedSystem for SAP HANA portfolio is comprised of the following SAP HANA certified appliances:

- **HPE ConvergedSystem 500 for SAP HANA:** Ideal for small to midsized SAP HANA deployments, these systems offer outstanding price/performance for scale-up analytics or S/4HANA workloads ranging from 128 GB of memory to scale-out workloads for data warehousing use cases with up to 68 TB of memory. Built with Intel® Xeon® E7 v4 architecture, the HPE ConvergedSystem 500 for SAP HANA is available in 2 CPU and 4 CPU configurations.
- **HPE ConvergedSystem 900 for SAP HANA:** Built with Intel Xeon E7 v4 architecture, and available in 2 CPU to 16 CPU configurations, the HPE ConvergedSystem 900 for SAP HANA is ideally suited for mission-critical S/4HANA use cases with up to 16 TB of memory.

Also available for SAP HANA Tailored Datacenter Integration use cases are TDI Compute Blocks based on the HPE ProLiant DL580 Gen9 server, which powers the HPE ConvergedSystem 500 for SAP HANA (CS500), and TDI Compute Blocks based on the HPE Superdome X, which powers the HPE ConvergedSystem 900 for SAP HANA (CS900).

The focus of this white paper is primarily on the SAP HANA Vora deployment with Hadoop as part of the total solution. Detailed information on the CS500 or CS900 is not included. However, the white paper will provide detail on integration with the CS500 infrastructure.

For a list of certified appliances from HPE for SAP HANA, refer to SAP’s online documentation at:

<https://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/appliances.html#categories=Hewlett%20Packard%20Enterprise%20Intel%20Broadwell%20EX%20E7>

More details on SAP HANA appliance are available in Appendix E.

HPE Elastic Platform for Big Data Analytics

The HPE Elastic Platform for Big Data Analytics (EPA) is a modular infrastructure foundation to accelerate business insights and enable organizations to rapidly deploy, efficiently scale and securely manage the explosive growth in volume, speed and variety of Big Data workloads. HPE supports two different deployment models under this platform:

- **HPE Balanced and Density Optimized (BDO) system** supports conventional Hadoop deployments that scale compute and storage together, with some flexibility in choice of memory, processor and storage capacity. This is primarily based on the HPE ProLiant DL380 server platform, with density optimized variants using the HPE Apollo 4200 and Apollo 4530 servers. Organizations that are already invested in balanced systems have the option of consolidating their existing deployments to a more elastic platform with the HPE Workload and Density Optimized (WDO) system.
- **HPE Workload and Density Optimized (WDO) system** harnesses the power of faster Ethernet networks to independently scale compute and storage using a building block approach, and lets you consolidate your data and workloads growing at different rates. The base HPE WDO system uses the HPE Apollo 4200 as a storage block and the HPE Apollo 2000 as a compute block. Additional building blocks such as HPE Moonshot, can be layered on top of the base configuration to target different workloads and requirements.

Figure 3 below highlights the different building blocks that are part of the HPE BDO and WDO system offerings. 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 building large-scale data lakes with diverse sets of data. As the workloads and data storage requirements change (often uncorrelated to each other) the HPE WDO system allows customers to easily scale by adding compute and storage blocks independently from each other, maximizing the infrastructure requirements for the workload demands.

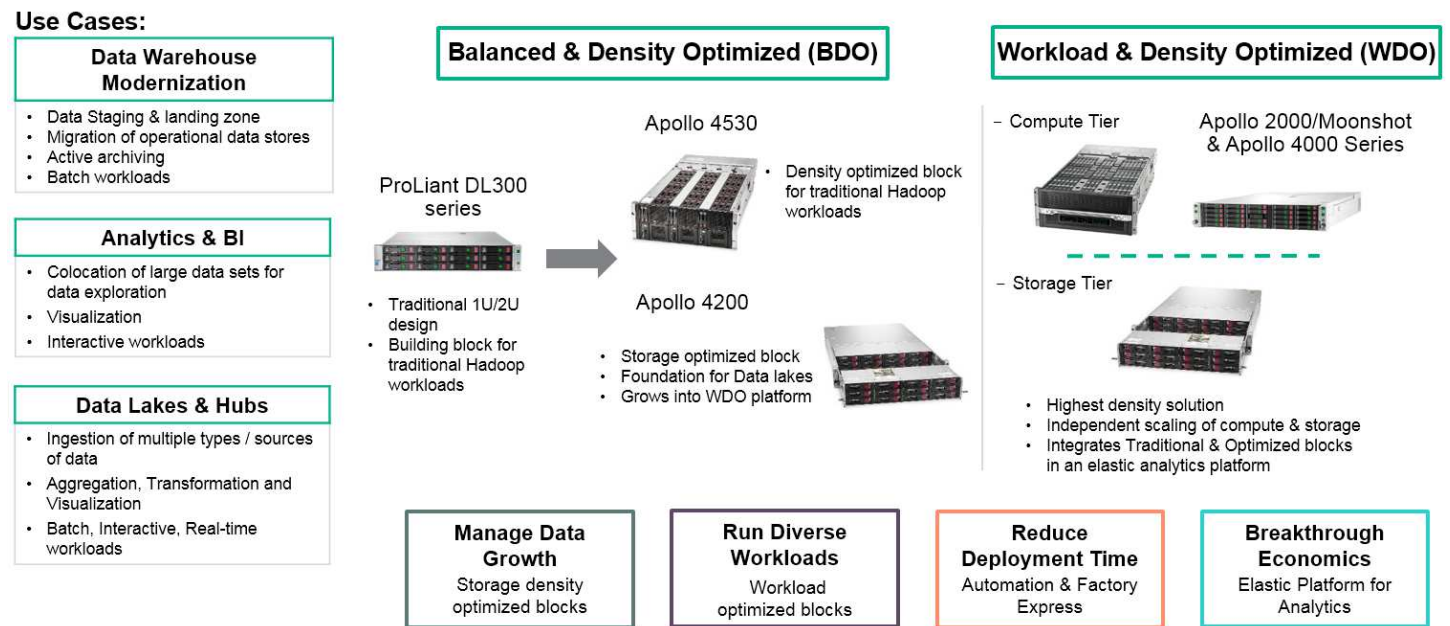


Figure 3. HPE Elastic Platform for Big Data Analytics with modular building blocks

Accelerators

An additional component of the HPE Elastic Platform for Big Data Analytics is the concept of Accelerators. Accelerators are specialized building blocks for optimizing workload performance, storage efficiency and deployment. As new workloads with more diverse requirements are added, accelerator building blocks are designed to target intended outcomes. Examples include accelerating performance of NoSQL databases such as HBase that require low-latency processing of events in near real time, in-memory analytics using Spark and SAP HANA Vora, deep learning on GPU accelerated servers, address storage efficiency with HDFS tiering and erasure coding, or deployment agility through automation and as-a-service solutions.

Figure 4 below provides an example multi-rack HPE WDO system supporting a wide variety of workloads leveraging a common HDFS data set. By separating out the compute and storage tiers, customers are able to address diverse workloads without having to duplicate the storage tier for each workload, which can be a costly, restrictive and siloed approach to solving disparate Big Data challenges.

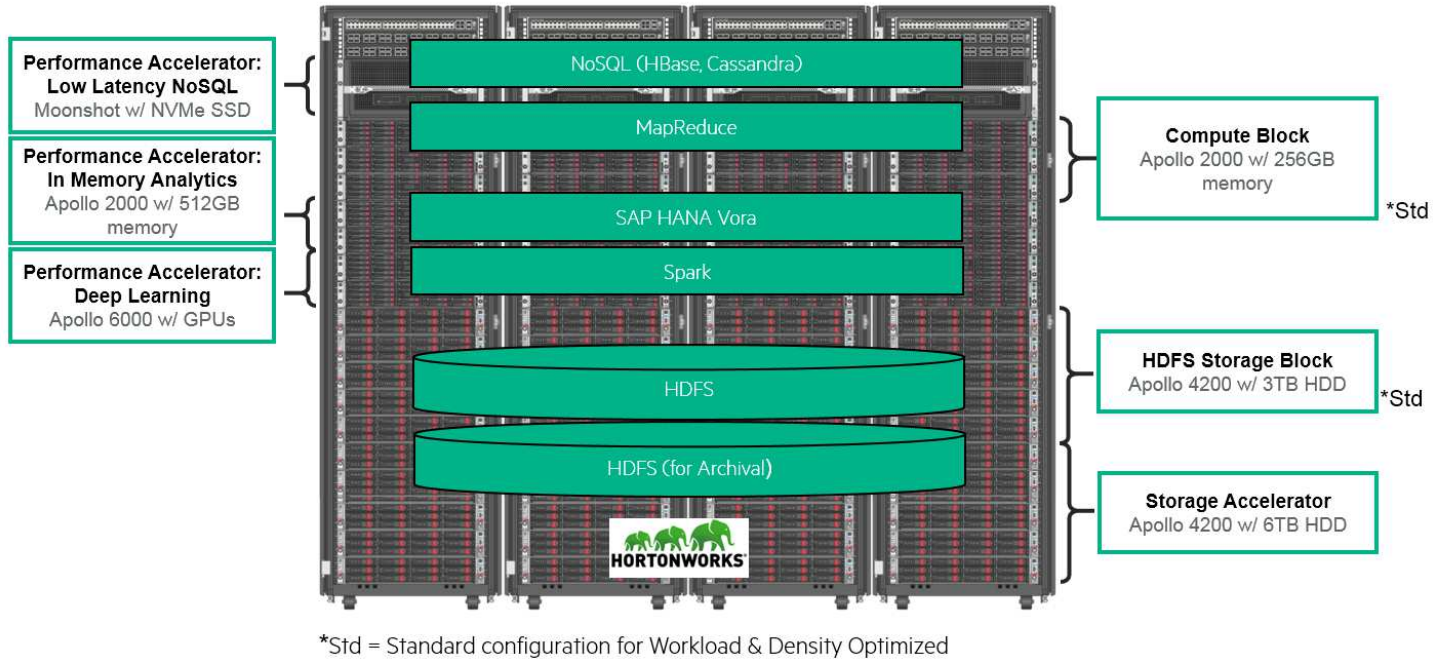


Figure 4. HPE EPA with Hortonworks HDP – Flexible architecture for Big Data workloads

The solution described in this white paper is composed of the HPE ConvergedSystem 500 for SAP HANA and HDP, and the HPE WDO system utilizing the HPE Apollo 2000 compute block to host the Spark and SAP HANA Vora services, and the HPE Apollo 4200 for data storage using HDFS (Hadoop distributed filesystem). By leveraging the HPE WDO system for the SAP HANA Vora solution, customers can ensure they have a scalable model for growing the compute and storage tiers independently as workload and data requirements change.

For a more in-depth analysis of the HPE ConvergedSystem for SAP HANA, and the various deployment choices available to deploy pre-configured appliances or tailored systems (TDI) that integrate with your existing data center infrastructure, refer to <http://h20195.www2.hpe.com/V2/GetDocument.aspx?docname=c04551209>

Software overview

Hortonworks Data Platform

The Hortonworks Data Platform (HDP), powered by Apache Hadoop, is a massively scalable and 100% open source platform for storing, processing and analyzing large volumes of data. It is designed to deal with data from many sources and formats in a very quick, easy and cost-effective manner.

HDP is a platform for multi-workload data processing across an array of processing methods – from batch through interactive and real-time – all supported with solutions for governance, integration, security and operations. As the only completely open Hadoop data platform available, HDP integrates with and augments your existing best-of-breed applications and systems so you can gain value from your enterprise Big Data, with minimal changes to your data architectures. Finally, HDP allows you to deploy Hadoop wherever you want it – from cloud or on-premises as an appliance, and across both Linux® and Microsoft® Windows®.

Hadoop 2.0 with YARN

Apache Hadoop is an open-source software framework for distributed storage and distributed processing of very large data sets, traditionally deployed on clusters of general purpose servers with local storage. All the modules in Hadoop are designed with a fundamental assumption that hardware failures are common and should be automatically handled by the framework.

The core of Apache Hadoop framework consists of:

- Hadoop Common: The common utilities that support the other Hadoop modules.
- Hadoop Distributed File System (HDFS): A distributed file system that provides high-throughput access to application data.
- Hadoop YARN: A framework for job scheduling and cluster resource management.
- Hadoop MapReduce: A YARN-based system for parallel processing of large data sets.

One of the most significant benefits of YARN and Hadoop 2.0, is to separate processing from resource management. Initial deployments of Hadoop focused on single batch-oriented workloads, and in order to run multiple workloads at the same time, multiple Hadoop clusters had to be stood up for each workload, resulting in cluster sprawl and data duplication.

With Hadoop 2 and YARN, new scheduling algorithms allow multiple workloads on the same cluster, accessing a common data pool along with hardware and software resources (see Figure 5). This enables a variety of new and familiar tools like Spark and SQL to identify data of value interactively and in real time, without being hampered by the often I/O intensive, high-latency MapReduce framework.

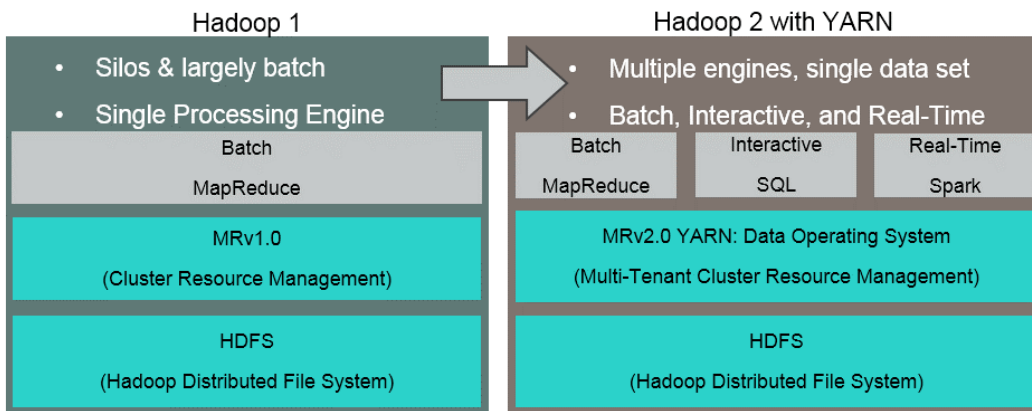


Figure 5. Evolution of Hadoop to enable multi-tenancy and diverse workloads with YARN

Apache Spark

Apache Spark is a powerful open source processing engine built around speed, ease of use, and sophisticated analytics. It was originally developed at UC Berkeley in 2009.

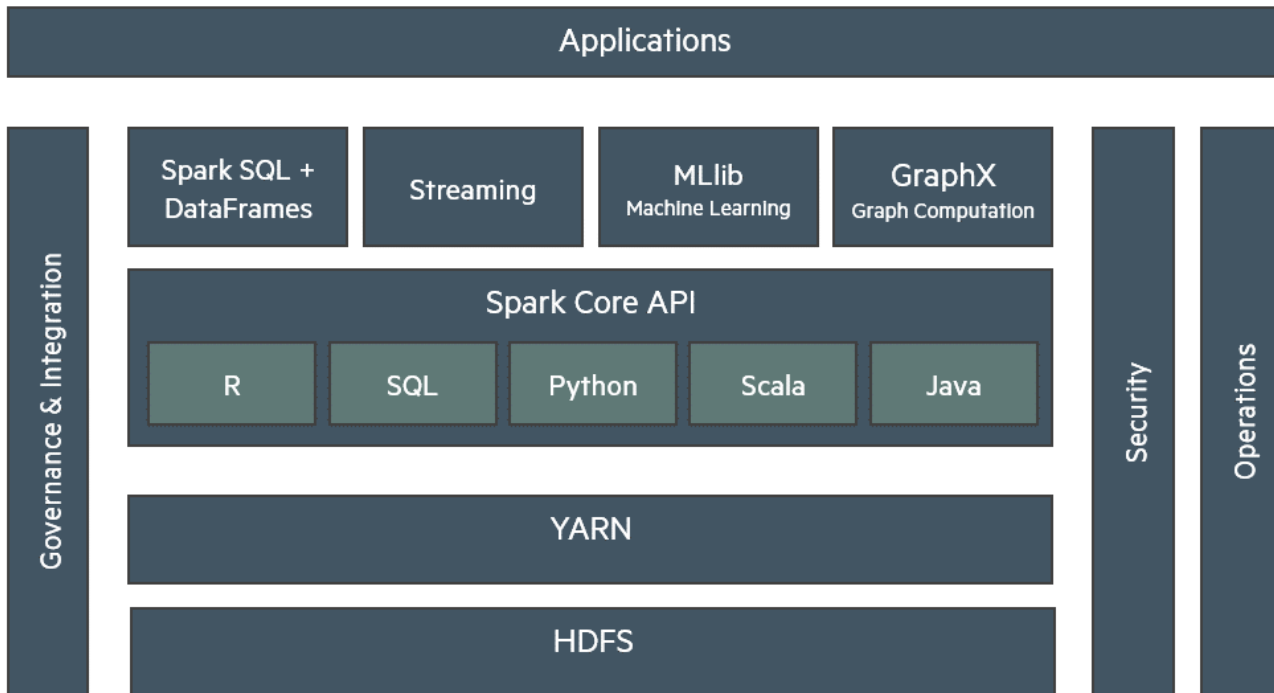


Figure 6. Apache Spark framework on HDP

Spark was developed in response to limitations in Hadoop’s two-stage disk-based MapReduce processing framework. Spark maintains MapReduce’s linear scalability and fault tolerance, and extends the processing capabilities in several important areas.

In-memory analytics

Instead of a traditional two-stage sequential map-then-reduce format which requires intermediate results to be written to disk, Spark enables in-memory access to intermediate results in a multi-stage processing pipeline through its Resilient Distributed Dataset (RDD) abstraction. This dramatically increases performance, in some cases up to 100 times faster than MapReduce.

Data federation

Spark has a rich set of libraries and APIs that enable developers to quickly build analytics workflows more efficiently to access any data source – from HDFS or Object storage to NoSQL and Relational databases, in batch or in real time. With Spark’s dataframe APIs, Hadoop data can now be integrated into Spark applications and workflows that extend beyond Hadoop and enable queries across relational data platforms such as SAP HANA using SAP HANA Vora, Vertica or Oracle.

Iterative analytics

Spark is well suited for highly iterative algorithms that are used for machine learning and graph analysis which require multiple passes while iteratively querying large in-memory data sets. Spark’s GraphX library unifies iterative graph analysis with ETL and interactive analysis.

Real-time streaming analytics

Spark also has strong integration with a variety of tools in the Hadoop ecosystem, and provides a unified platform that can be used to process streams of data from Flume or Kafka in micro-batches using Spark Streaming. With Spark, developers can now run a single codebase that can be used for both real-time streaming data as well as batch analytics.

SAP HANA

SAP HANA is the foundation of SAP’s digital enterprise platform, to simplify data management and drive processing agility for SAP’s connected business applications by providing a single multi-tenant platform with application building blocks for database, processing, integration and application services. SAP HANA, short for “High-performance Analytic Appliance,” is an in-memory database and application platform that offers significant performance benefits over conventional database platforms for both Online Analytical Processing (OLAP) and Online Transaction Processing (OLTP). SAP HANA systems can scale up or scale out to handle in-memory processing of terabytes of data

SAP HANA has now introduced the concept of data temperatures and tiering to manage the cost of data storage and processing at the database storage layer. This extends the platform to intelligently distribute data and its processing to drive economies of scale by moving warm and cold data off memory to more economical disk based alternatives like Hadoop.

SAP HANA Vora

SAP HANA Vora is a critical component of SAP’s Digital Enterprise Platform that provides contextually aware analytics capability by integrating the SAP HANA platform seamlessly with Hadoop. SAP HANA Vora is an in-memory query engine that brings powerful contextual analytics across all data stored in Hadoop, enterprise systems and other distributed data sources, and drives lower TCO by achieving low cost yet high speed analytics on data volumes that are too large to be processed by SAP HANA alone. It also extends the SAP HANA platform with the ability to store Big Data and has data temperature management to facilitate data movement between SAP HANA, SAP HANA Vora and Hadoop.

SAP HANA Vora extends Spark with a richer SQL capability that supports the data hierarchies needed for analytics, and optimizes Spark SQL query execution on HDFS data. SAP HANA Vora accelerates query performance through a high-speed caching layer on top of Spark, and uses a SQL interpreter that uses C libraries, and distributes the query processing in parallel on each node of a Spark cluster. To enable the use of the SAP HANA Vora processing engine and its functional features in applications, SAP HANA Vora provides enhanced data source API implementations and Spark extensions.

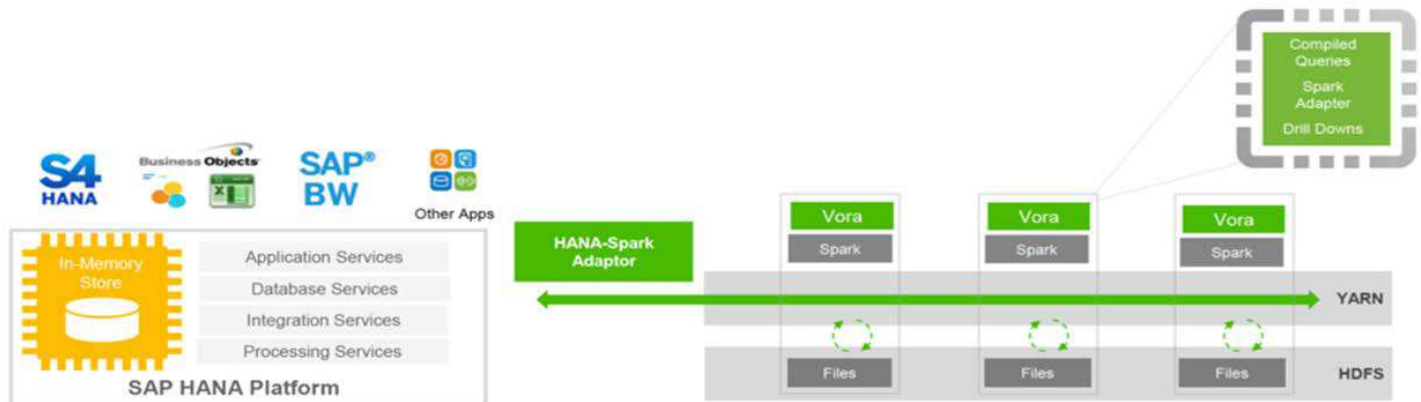


Figure 7. SAP HANA Vora Overview (Source: SAP)

SAP HANA Vora acts as a bridge between SAP HANA and Hadoop, and enables several key business use cases as described in the following section.

SAP HANA Vora use cases

Data tiering

The most basic use case for existing SAP HANA customers is the ability to use Hadoop as a cold data store for less frequently accessed data that can be moved from SAP Business Suite or S/4HANA using SAP's data tiering capabilities. This could include mostly read-only data like sales documents, Customer Case notes attachments, etc. As Figure 8 below shows, data tiering allows enterprises to transparently manage large data volumes by automatically moving data among memory, disk and Hadoop/SAP IQ. The tool that enables it is the Data Lifecycle Manager (DLM), a component of SAP HANA's data warehouse foundation.

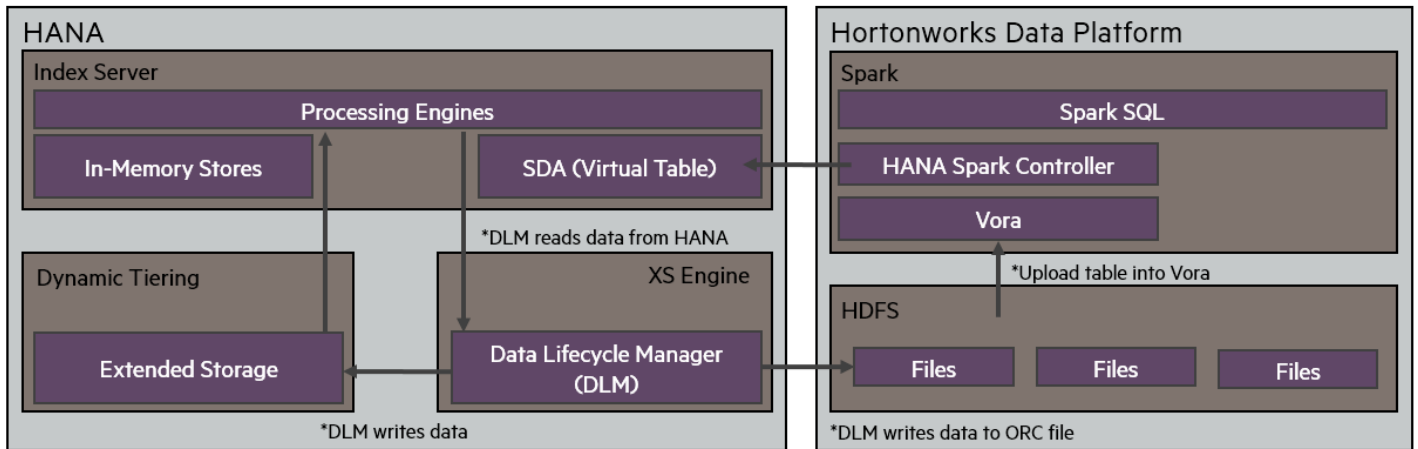


Figure 8. SAP HANA Vora use case – data tiering (Source: SAP)

Analytics

While data tiering allows for bi-directional data movement, the key driver for SAP HANA Vora is the ability to interactively analyze data with a single logical view that ties business data in SAP HANA with Big Data, while storing it in two distinct physical systems. This provides data scientists access to real datasets across both sources without having to move data.

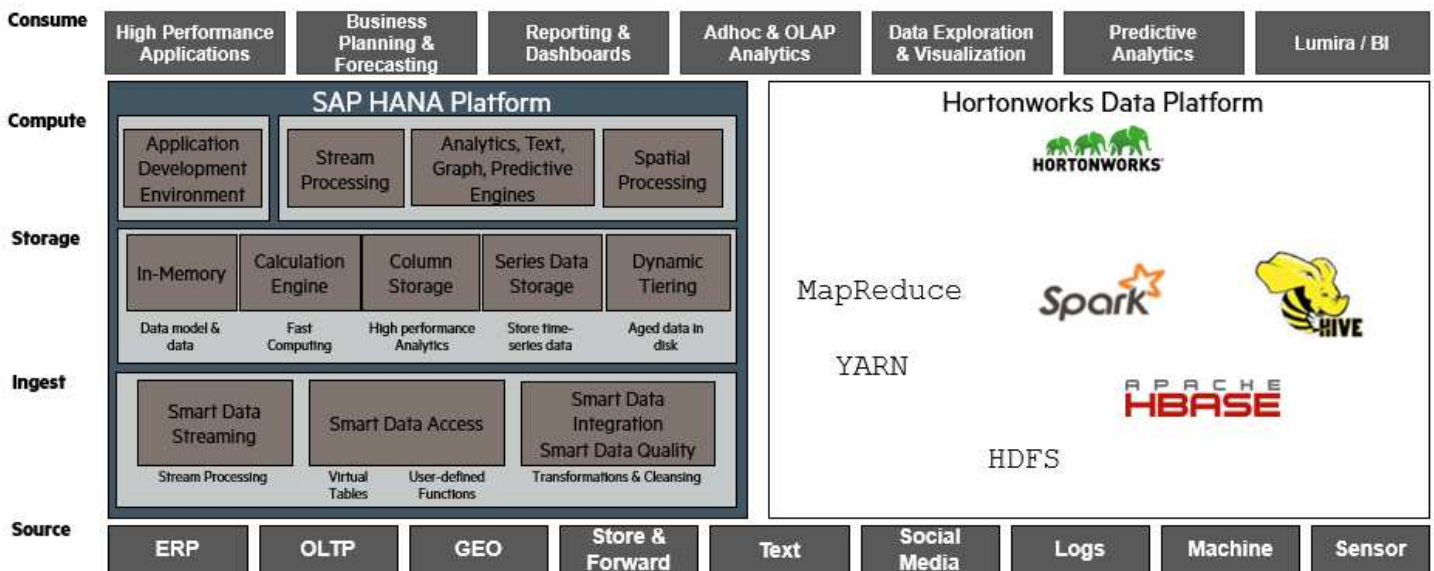


Figure 9. SAP HANA Vora use case – analytics (Source: SAP)

As shown in Figure 9 above, the integration of SAP HANA Vora with SAP HANA and HDP drives new levels of performance for analytics use cases that include but are not limited to

- Internet of Things (IoT) – Preventive maintenance by analyzing large volumes of sensor data in real time
- Manufacturing – analyzing manufacturing test data to drive better real-time packaging decisions
- Telecom – network capacity planning and traffic shaping for the telecom industry
- Finance – fraud detection and risk mitigation for the financial sector

Solution components

Utilizing HPE EPA/WDO and HPE ConvergedSystem 500, Figure 10 provides a basic conceptual diagram of SAP HANA Vora on the HPE WDO system.

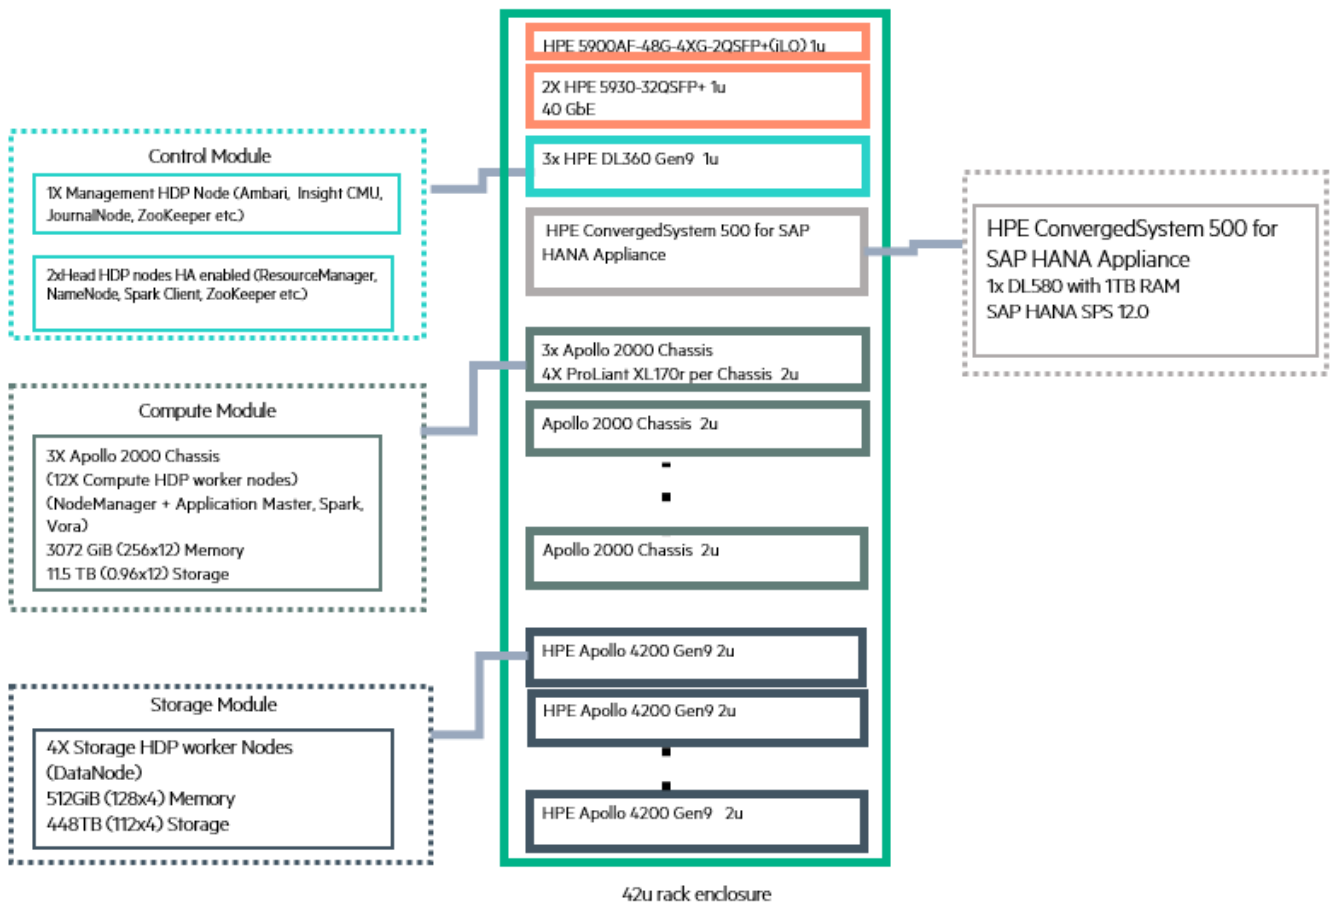


Figure 10. Integrated solution – SAP HANA Vora on HPE EPA with HDP and CS500 for SAP HANA appliance

For full BOM listings of products selected for the proof-of-concept, refer to [Appendix A](#) of this white paper.

Hardware components

Figure 11 shows the reference hardware details of the solution, with 12 worker nodes and 4 storage nodes housed in a single 42U rack along with the HPE CS500 appliance.

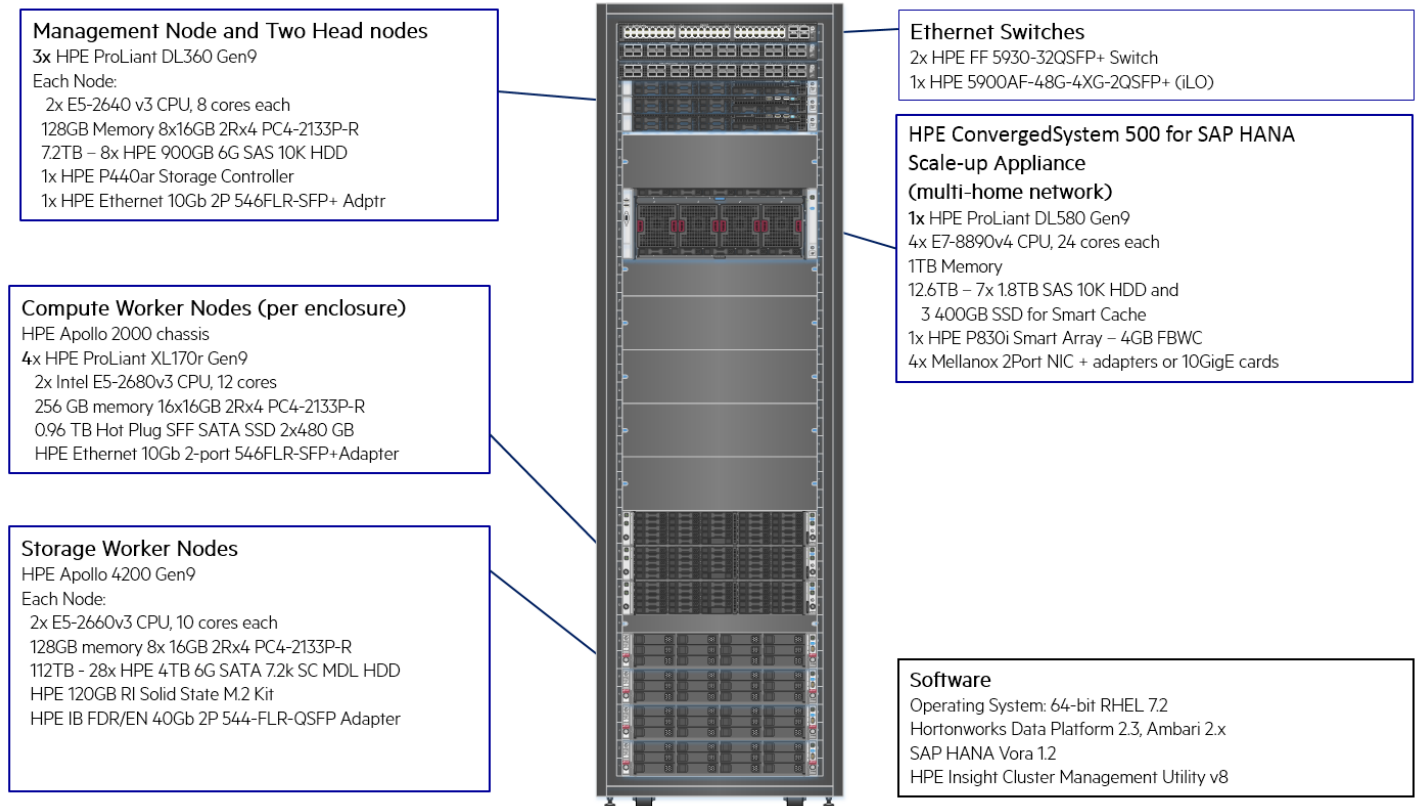


Figure 11. Base HPE EPA configuration with HPE ConvergedSystem 500 for SAP HANA appliance with lists of common components

HPE ConvergedSystem for SAP HANA

HPE ConvergedSystem 500 for SAP HANA Scale-up and Scale-out configurations are certified SAP HANA appliances. These are preconfigured SAP HANA appliances that act as database servers with more than a 1000 combinations of options, including memory size, number of processors, type of processors, disk configurations, operating system, and encryption option. These are Intel Xeon E7 v4 processor-based servers with SAP HANA version SPS 12. For this paper, the HPE CS500 configuration consisted of 1x HPE ProLiant DL580 Gen9 with 4x E7-8890 v4 CPU with 24 cores, 1 TB memory, 12.6 TB on SAS 10k HDDs, HPE Smart Array P830i and 4 10 Gb Ethernet cards.



Figure 12. HPE ConvergedSystem 500 for SAP HANA Scale-up configuration

HPE Workload and Density Optimized (WDO) system building blocks

This Reference Architecture uses the building blocks of the HPE Workload and Density Optimized (WDO) system to deploy HDP, Spark and SAP HANA Vora.

Compute nodes – HPE Apollo 2000

The reference HPE WDO system configuration features three HPE Apollo 2000 chassis containing a total of 12 HPE ProLiant XL170r Gen9 hot-pluggable server nodes. The HPE ProLiant XL170r Gen9 is shown in Figure 13. The rear and front of the HPE Apollo 2000 chassis is shown in Figure 14.



Figure 13. HPE ProLiant XL170r Gen9



Figure 14. HPE Apollo 2000 chassis (rear and front)

The HPE Apollo 2000 system is a dense solution with four independent HPE ProLiant XL170r Gen9 hot-pluggable server nodes in a standard 2U chassis. Each HPE ProLiant XL170r Gen9 server node is serviced individually without impacting the operation of other nodes sharing the same chassis to provide increased server uptime. Each server node harnesses the performance of up to 2400 MHz memory (16 DIMM slots per node) and dual Intel Xeon E5-2600 v3 or v4 series processors in a very efficient solution that shares both power and cooling infrastructure. Other features of the HPE ProLiant XL170r Gen9 server include:

- Support for high-performance memory (DDR4) and Intel Xeon E5-2600 v3 and v4 processor up to 22C, 145W
- Additional PCIe riser options for flexible and balanced I/O configurations
- FlexibleLOM feature for additional network expansion options
- Support for dual M.2 drives

For more information on HPE Apollo 2000 chassis, visit hpe.com/us/en/product-catalog/servers/proliant-servers/pip.hpe-apollo-r2000-chassis.7832023.html

For more information on the HPE ProLiant XL170r Gen9 server, visit hpe.com/servers/xl170r

Each of these compute nodes typically runs HDP's YARN, Spark and SAP HANA Vora services.

Storage nodes – HPE Apollo 4200

The reference HPE WDO system configuration features four HPE Apollo 4200 Gen9 servers. Each server is configured with 28 LFF disks; each typically runs HDFS DataNode. The HPE Apollo 4200 Gen9 server is shown in Figure 15.



Figure 15. HPE Apollo 4200 Gen9 server

The HPE Apollo 4200 Gen9 server offers revolutionary storage density in a 2U form factor for data intensive workloads such as Apache Spark on Hadoop HDFS. The HPE Apollo 4200 Gen9 server allows you to save valuable data center space through its unique density optimized 2U form factor which holds up to 28 LFF disks and with capacity for up to 224 TB per server. It has the ability to grow your Big Data solutions with its density optimized infrastructure that is ready to scale. Another benefit is that the HPE Apollo 4200 fits easily into standard racks with a depth of 32-inches per server – no special racks are required.

For more detailed information: hpe.com/us/en/product-catalog/servers/proliant-servers/pip.hpe-apollo-4200-gen9-server.8261831.html

Key point

The storage controllers in the HPE Apollo 4200 and HPE ProLiant XL170r support HPE Secure Encryption. Secure Encryption is an HPE Smart Array controller-based data encryption solution. It provides encryption for data at rest, an important component for complying with government regulations having data privacy requirements, such as HIPAA and Sarbanes-Oxley. This optional feature enhances the functionality of the storage controller in cases where encryption is a required feature of the solution.

Management/head nodes

Three HPE ProLiant DL360 Gen9 servers are configured as management/head nodes. Refer to Table 2 for the suggested software components for these nodes.

Management and head nodes store important metadata information for the cluster and hence for greater data protection and performance a RAID configuration is suggested for data storage. RAID 1 is recommended for the OS and NameNode metadata. RAID 10 is recommended for database data.

Edge node – The HPE WDO system includes an optional multi-homed network server called Edge node. It acts as a gateway node between the cluster's private VLAN and the external routable network. Any application that requires both external network access and cluster-private-network can run on this server. When significant storage and network bandwidth are required, use additional HPE ProLiant DL360 Gen9 servers as Edge nodes.

Power and cooling

When planning large clusters, it is important to properly manage power redundancy and distribution. To ensure servers and racks have adequate power redundancy, HPE recommends that each chassis (HPE Apollo 2000) and each HPE Apollo 4200 Gen9 server should have a backup power supply and each rack should have at least two Power Distribution Units (PDUs).

There is additional cost associated with procuring redundant power supplies; however, the need for redundancy is less critical in larger clusters where the inherent redundancy within HDP ensures there would be less impact in the event of a failure.

Best practice

For each chassis, HPE Apollo 2000, Apollo 4200 Gen9, and HPE ProLiant DL360 Gen9 server, HPE recommends connecting each of the device’s power supplies to a different PDU. Furthermore, PDUs should each be connected to a separate data center power line to protect the infrastructure from a line failure.

Distributing server power supply connections evenly to the PDUs while also distributing PDU connections evenly to data center power lines ensures an even power distribution in the data center and avoids overloading any single power line. When designing a cluster, check the maximum power and cooling that the data center can supply to each rack and ensure that the rack does not require more power and cooling than is available.

Networking

As shown in Figure 16, two IRF-bonded HPE FlexFabric 5930-32QSFP+ switches are specified in each rack for high performance and redundancy. Each provides six 40 GbE uplinks that can be used to connect to the desired network, or, in a multi-rack configuration, to another pair of HPE FlexFabric 5930-32QSFP+ switches that are used for aggregation.

For the networking requirements of Hadoop and SAP HANA Vora, some ports on the HPE FlexFabric 5930 switch are set aside for storage, compute, SAP HANA Vora and other networks. Ports 1-8 are assigned for storage nodes; ports 9-17 are assigned for compute nodes using QSFP+ 4x10G SFP+ 3m DAC cables, splitting 40G into 4x10G ports; ports 18-19 are used to connect to SAP HANA Vora server or switch depending on the requirements. Ports 21-22 are used for connecting HPE ProLiant DL360 Gen9 Hadoop master nodes. Ports 23-28 can be used for customer uplinks, and ports 29-32 are used for IRF.

In the reference configuration, ports 1-4 are connected to HPE Apollo 4200 storage worker nodes with a 40 GbE network. Any future expansion of storage nodes will use ports above 4 in ascending order. Ports 15-17 are connect to HPE Apollo 2000 compute worker nodes with a 10 GbE network using QSFP+ 4x10G SFP+ 3m DAC cables. Any future expansion of compute nodes will use the ports below 15 in descending order. Port 19 is connected to the SAP HANA Vora server with a 40 GbE network. Port 21 is connected to HPE ProLiant DL360 Gen9 management and head nodes using QSFP+ 4x10G SFP+ 3m DAC cables with a 10 GbE network. All these networks are using two bonded NICs on each node.

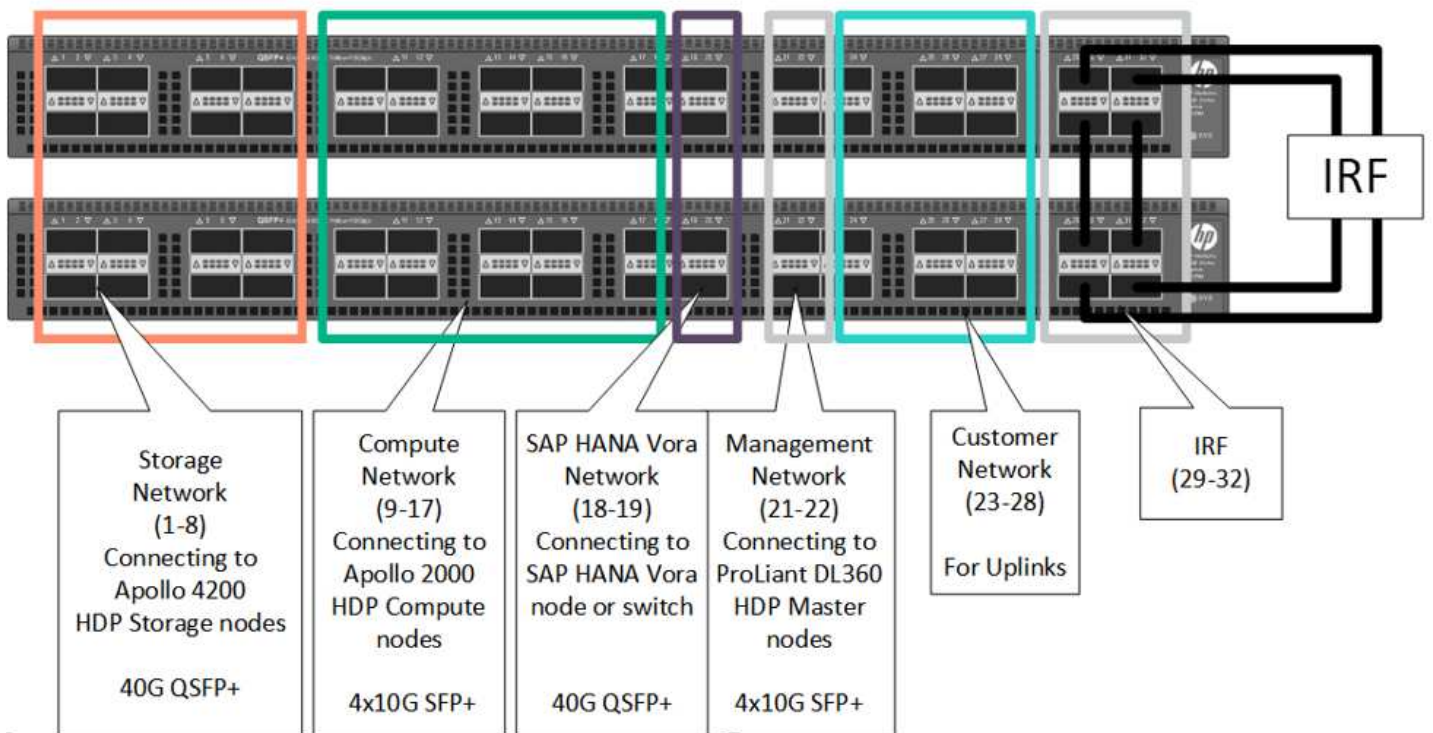


Figure 16. HPE EPA networking

Note

IRF bonding requires four 40 GbE ports per switch, leaving six 40 GbE ports on each switch for uplinks. QSFP+ 4x10 G SFP+ DAC cables are used to connect compute and management nodes to the QSFP+ ports on the HPE FlexFabric 5930 switches. Therefore a single QSFP+ port on the HPE FlexFabric 5930 switch handles 4 compute node connections.

iLO network

A single HPE FlexFabric 5900 switch is used exclusively to provide connectivity to HPE Integrated Lights-Out (iLO) management ports, which run at or below 1 GbE. The iLO network is used for system provisioning and maintenance.

Cluster isolation and access configuration

It is important to isolate the Hadoop cluster so that external network traffic does not affect the performance of the cluster. In addition, isolation allows the Hadoop cluster to be managed independently from its users, ensuring that the cluster administrator is the only person able to make changes to the cluster configuration.

Thus, HPE recommends deploying ResourceManager, NameNode, and Worker nodes on their own private Hadoop cluster subnet.

Key point

Once an HDP / Apache Hadoop cluster is isolated, the users still need a way to access the cluster and submit jobs to it. To achieve this, you can use edge nodes that have access to both the Hadoop cluster subnet and a subnet belonging to users. Typically, edge nodes in production clusters are dedicated for client-side operations of a Hadoop cluster, and provide access to the Hadoop cluster. For smaller clusters, these nodes are shared by multiple services (Resource Manager, NameNode, etc.).

In cases where additional network connections are required to speed up large scale data ingestion for either batch or real-time data processing, using additional edge nodes and employing an externally addressable, routable subnet for compute and storage nodes is recommended.

HPE recommends installing client gateways on the multi-homed edge node so that users can shell into the edge node and submit jobs to the cluster.

Data ingestion

After the HDP / Apache Hadoop cluster has been isolated on its own private network, you must determine how to access HDFS in order to ingest data. The HDFS client must be able to reach every Hadoop data node in the cluster in order to stream blocks of data on to the HDFS.

HPE provides the following options for staging data:

- **Using an edge node** – A multi-homed edge node can be used to stage data. HPE recommends configuring this server with eight disks to provide a sufficient amount of disk capacity to provide a staging area for ingesting data into the Hadoop cluster from another subnet.
- **Using the ToR switches** – You can make use of open ports in the ToR switches. This is designed such that if both NICs on each storage node are used, eight ports are used on each HPE Apollo 2000 chassis (two NICs on each node), and two NICs are used on each management node, the remaining 40 GbE ports on the ToR switches can be used by multi-homed systems outside the Hadoop cluster to move data into the cluster.

Software components**Hortonworks Data Platform**

Hortonworks Data Platform (HDP) is a platform for multi-workload data processing. HDP can utilize a range of processing methods – from batch to interactive and real time – all supported by solutions for governance, integration, security, and operations. HDP integrates with and augments solutions like HPE WDO system, allowing you to maximize the value of Big Data. HDP enables Open Enterprise Hadoop, a full suite of essential Hadoop capabilities in the following functional areas: data management, data access, data governance and integration, security, and operations.

HDP provides two key groups of services – Management Services and Worker Services. Management Services manage the cluster and coordinate the jobs whereas Worker Services are responsible for the actual execution of work on the individual scale out nodes. The RA provided in this document maps the Management and Worker services onto HPE infrastructure for clusters of varying sizes.

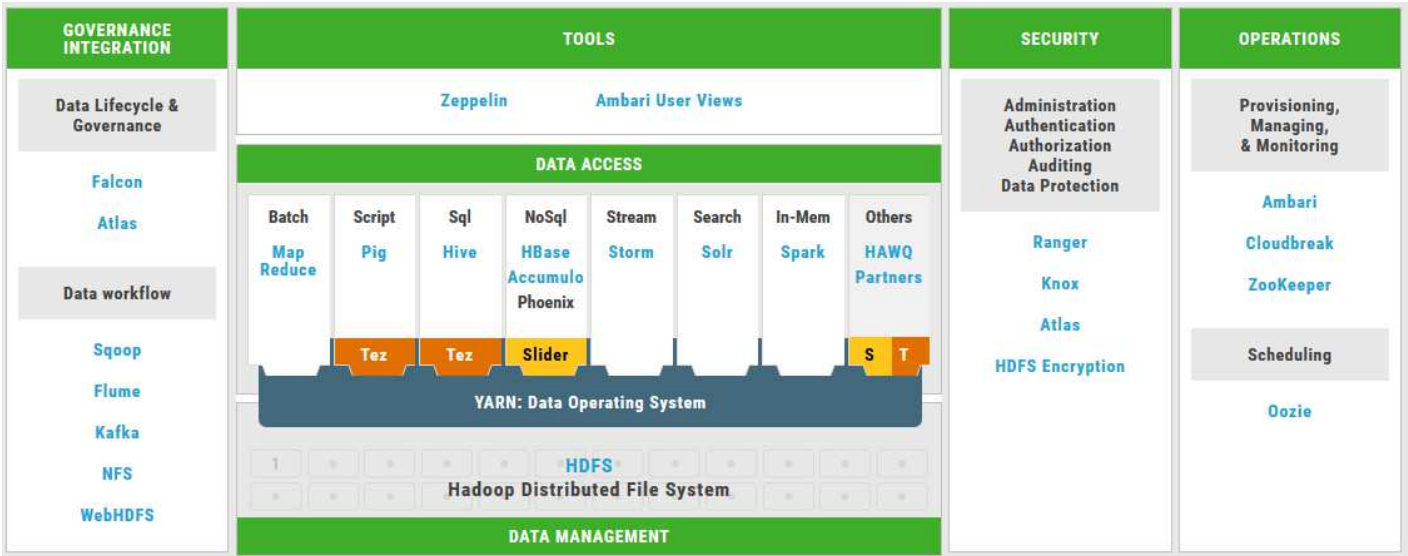
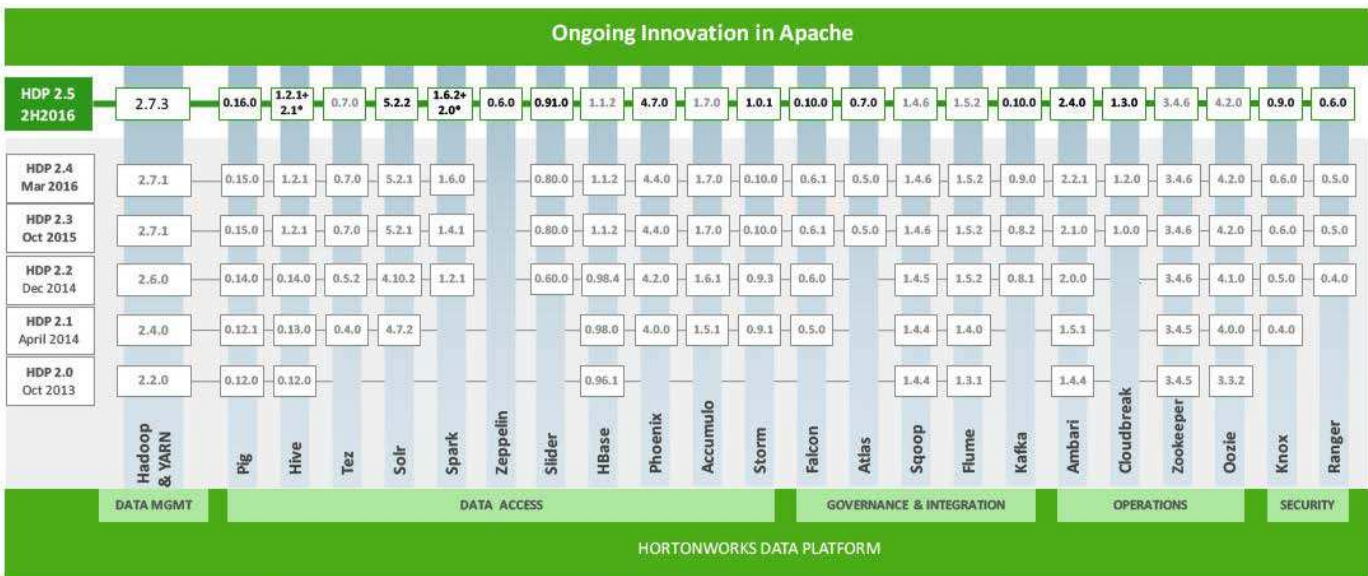


Figure 17. Hortonworks Data Platform: A full enterprise Hadoop data platform

YARN and Hadoop Distributed File System (HDFS) are the cornerstone components of HDP. While HDFS provides the scalable, fault-tolerant, cost-efficient storage for your Big Data lake, YARN provides the centralized architecture that enables you to process multiple workloads simultaneously. YARN provides the resource management and pluggable architecture for enabling a wide variety of data access methods.

Note: The reference solution tested used the HDP 2.3 distribution from Hortonworks, the HDP release compatible with SAP HANA Vora 1.2. Since then, Hortonworks announced the availability of HDP version 2.5.

Figure 18 shows all the Hadoop ecosystem components and their versions that are packaged in HDP releases up to the latest release, HDP 2.5. For detailed information on Hortonworks Data Platform, see hortonworks.com/hdp



* Spark 1.6.2+ Spark 2.0 – HDP 2.5 support installation of both Spark 1.6.2 and Spark 2.0. Spark 2.0 is Technical Preview within HDP 2.5. Hive 1.2.1+ Hive 2.1 – Hive 2.1 is Technical Preview within HDP 2.5.

Figure 18. HDP 2.5 ecosystem component details

Apache Spark

Apache Spark is an open source data processing engine built for speed, ease of use and sophisticated analytics. In recent years, Apache Spark has captured the Big Data computing market by providing speed and other improvements that are highly significant in comparison to MapReduce. Spark offers batch, interactive and streaming computational models. It provides a unified framework to manage Big Data processing needs with a variety of datasets that are varied in nature, such as text, graph data, time-series from batch and real-time streaming data sources.

Note: HDP 2.3 distribution includes Apache Spark 1.3.1, but SAP HANA Vora requires Apache Spark 1.5.2. Hence, download Apache Spark 1.5.2 from <http://spark.apache.org>

Apache Zeppelin

Apache Zeppelin is a web-based notebook tool that enables interactive data analytics. It allows you to make beautiful data-driven, interactive and collaborative documents with SQL, Scala, SAP HANA Vora and more. The Apache Zeppelin interpreter concept allows any language/data-processing-backend to be plugged into Zeppelin. Especially, Apache Zeppelin provides built-in Apache Spark integration. SAP HANA Vora provides an extension that supports SAP HANA Vora SQL interpreter.

SAP HANA Vora requires Zeppelin 0.5.6 built against Spark 1.5.2, Hadoop 2.6, and YARN.

SAP HANA

The SAP HANA Enterprise and Platform Editions provide the tools and services required for integration with Hadoop through SAP HANA Vora and SAP DLM. SAP HANA has three major components: Application services, database services and integration services (including developer APIs).

The integration services component provides the integration between SAP HANA and external systems like Spark and Hadoop, and includes features like Smart Data Access for data virtualization and federation, bulk loading, bi-directional real-time replication, data transformation, data cleansing and streaming data processing capabilities. Remote data synchronization enables companies to sync information between enterprise and remote devices and locations as well as Hadoop integration with Hive, Spark and direct access to HDFS files and MapReduce jobs using virtual user-defined functions (vUDF).

SAP HANA Vora

SAP HANA Vora is a distributed in-memory processing framework that extends Spark with a richer SQL capability that supports the data hierarchies needed for analytics, and optimizes Spark SQL query execution on HDFS data. SAP HANA Vora accelerates query performance through a high-speed caching layer on top of Spark, and distributes the query processing in parallel on each node of a Spark cluster. SAP HANA Vora enables integration of SAP HANA and Hadoop environments and data movement between SAP HANA and Hadoop. It provides enhanced data source API implementations that enable Spark SQL queries or parts of the queries to be pushed down to the SAP HANA Vora processing engine and also makes OLAP-style capabilities available in SQL.

SAP HANA DLM

Data Lifecycle Manager is a generic database-driven tool that enables you to model aging rules on SAP HANA tables in order to relocate "aged" or less frequently used data from SAP HANA tables in native SAP HANA applications. With Data Lifecycle Manager, SAP HANA Data Warehousing Foundation provides an SAP HANA XS-based tool to relocate data in native SAP HANA use cases from SAP HANA persistency to storage locations such as the SAP HANA Dynamic Tiering option, Hadoop (Spark SQL) or SAP IQ.

A proof-of-concept for interactive analytics and data tiering with SAP HANA Vora

The solution developed in this document focuses on the implementation and functionality validation of two key features of SAP HANA Vora:

- Interactive analytics across both SAP HANA and Hadoop data using SAP HANA Vora and Spark
- Manage data movement between SAP HANA and the Hortonworks Data Platform powered by Apache Hadoop

Interactive analytics across SAP HANA and Hadoop

Using a predictive analytics use case for preventive maintenance, we can identify part failures due to high temperature, check parts availability and proactively replace faulty parts. Machine part data is maintained in the SAP HANA database, while raw machine sensor data is collected from remote sites and stored in Hadoop (Figure 19).

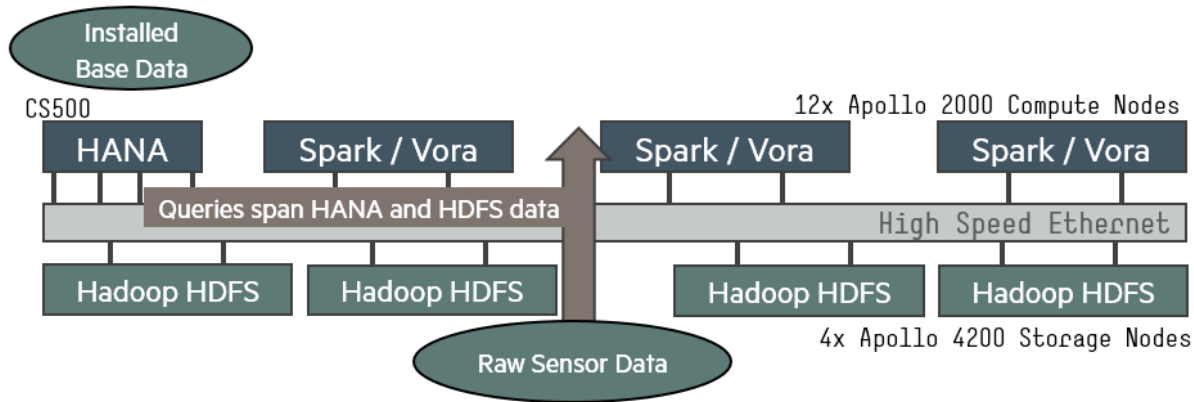


Figure 19. Federating interactive queries between SAP HANA and HDP powered by Apache Hadoop

Our sample company has several plants around the world with sensors attached to various machine parts. The SAP HANA database stores the installed base data of all machines and parts including the plant name and location, part age and cost, and available quantity.

The sensors attached to the machine parts are monitoring their temperature and generate data about collection times, dates and observed temperature. This data can be typically streamed through one of Spark, Storm or Flume and stored in HDFS, before being transferred into SAP HANA Vora. For our tests, the data was manually ingested into Hadoop.

The test configuration consisted of 100 GB of sensor temperature data loaded into SAP HANA Vora memory tables and 20 GB of machine parts data loaded into SAP HANA memory tables. The SAP HANA Vora tables are loaded from HDFS using VORA CREATE TABLE statement using Zeppelin Vora interpreter.

We can now run interactive queries and visualize the results using Zeppelin Vora Interpreter and Zeppelin visualization features across machine part data in SAP HANA tables and sensor data in SAP HANA Vora tables.

In order to analyze the number of machine parts that are running “hot”, i.e., exceeding the target operating temperature by 7 degrees or more, and sort by location, using Zeppelin, we can run an interactive query against the SAP HANA table HANA_MACHINE_PARTS and the SAP HANA Vora table vora_sensortemps, as shown in Figure 20.

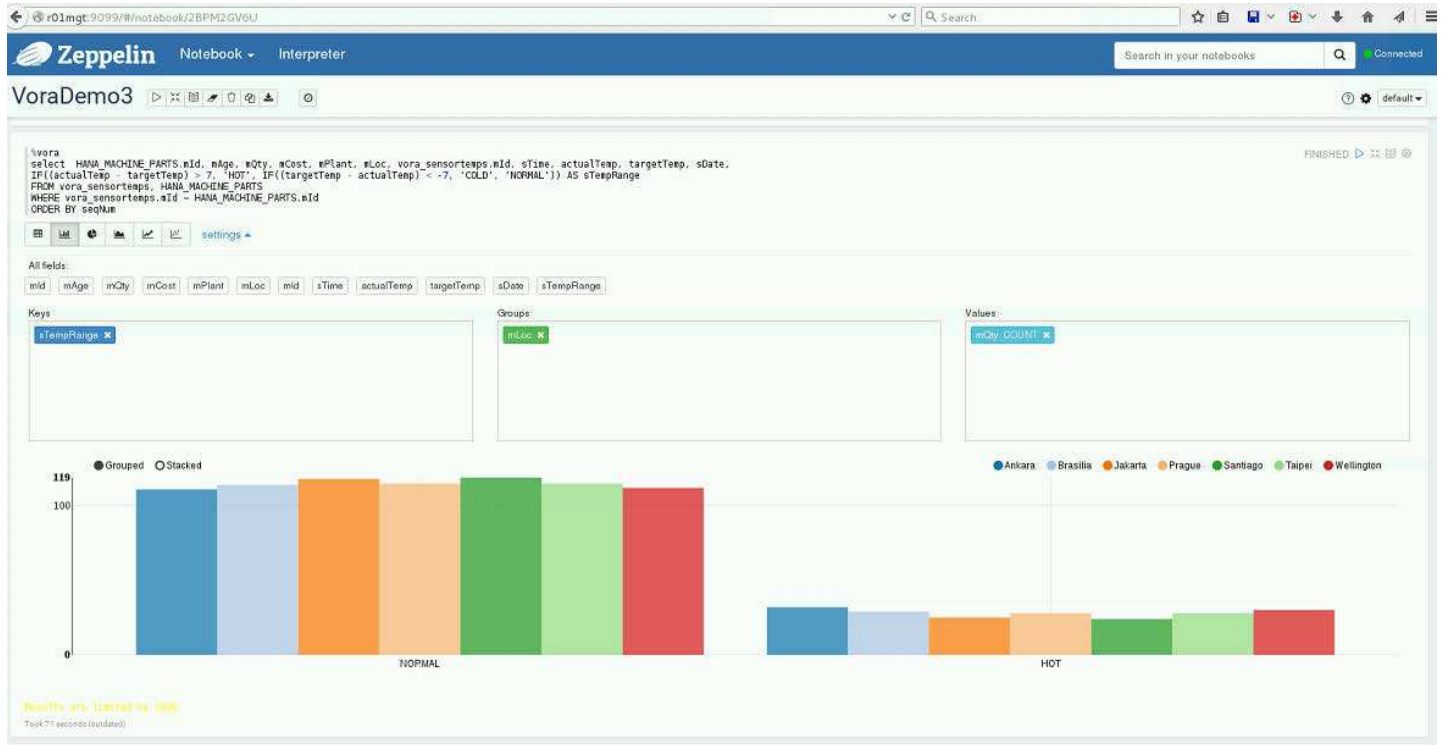


Figure 20. A sample federation query using Zeppelin

In order to calculate the dollar amount, by location for budgetary purposes, to replace the machine parts that are running hot, we can run an aggregation query joining the machine part cost field in an SAP HANA table and the sensor temperature data in an SAP HANA Vora table, as shown in Figure 21.

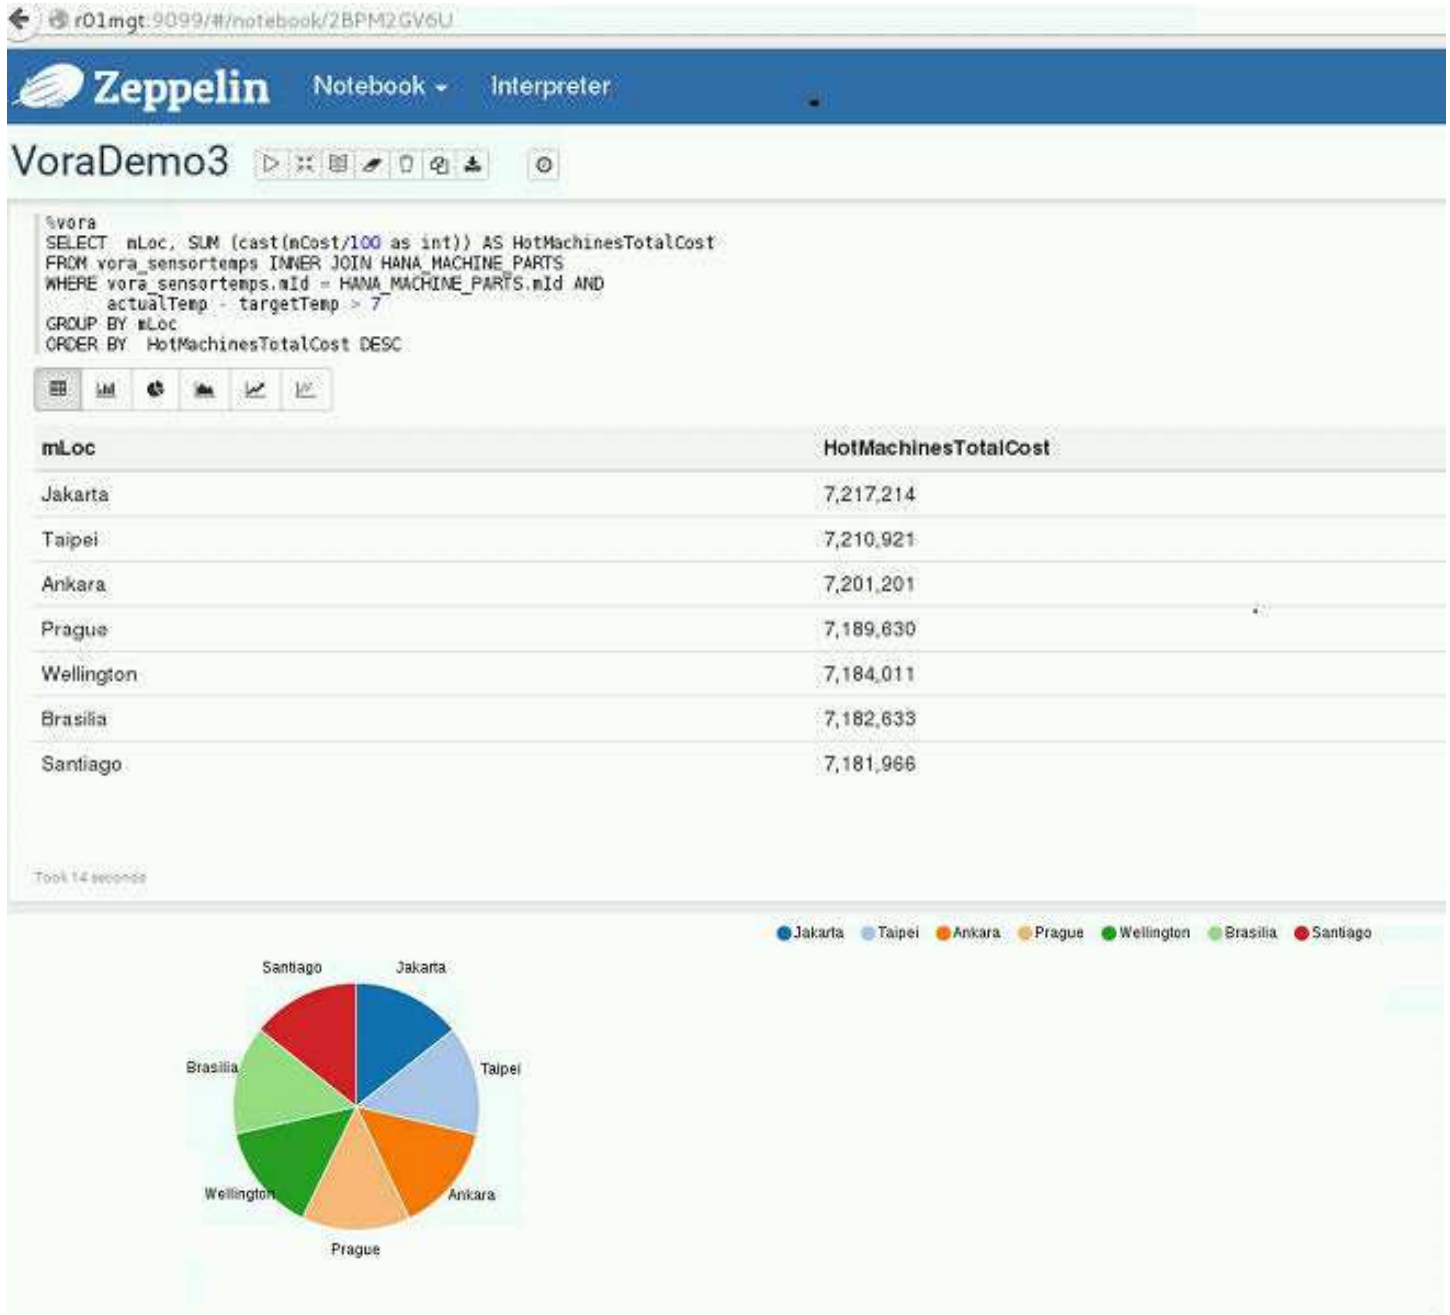


Figure 21. A sample aggregation query using Zeppelin

This is just one example of an interactive query that provides a logical view of SAP HANA and Hadoop data without physically moving any data. SAP HANA with SAP HANA Vora also enables OLAP analytics for HDP / Apache Hadoop data using hierarchical data structures that allow for complex computations on different levels of data. Extensions to Spark SQL also include enhancements to the data source API to enable Spark SQL queries or parts of the queries to be pushed down to the SAP HANA Vora processing engine.

Data movement between SAP HANA and HDP / Apache Hadoop

There are multiple options available for data movement in either direction between SAP HANA and Hadoop. SAP DLM provides both a tool-based approach to data lifecycle management, and automation with scheduling for mass data movement. The implementation details for these approaches are beyond the scope of this document. However, we were able to validate basic data movement functionality by running a script to push data from SAP HANA to HDP / Apache Hadoop HDFS using SAP HANA Vora and Spark SQL, as shown in Figure 22. The data was stored in HDFS in Parquet and ORC formats which enable efficient processing by SAP HANA Vora, Spark SQL and other processing frameworks. In actual deployments, this data movement is accomplished through configurations in SAP HANA DLM in a more automated and transparent manner.

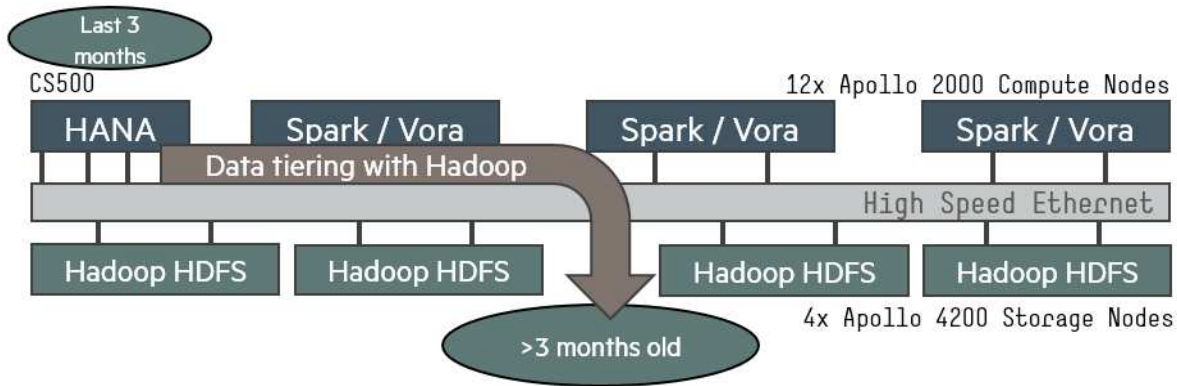


Figure 22. Data movement from SAP HANA to HDP powered by Apache Hadoop

Solution configuration guidance

The configurations recommended below are based on lab test scenarios and as such configuration settings mentioned in this section are intended to assist in optimizing the setup for the various services of this proof of concept. They are recommended to be used as a starting point for actual deployment scenarios. Customizations will likely be required in real life deployments.

The instructions provided here assume you have already created your Hadoop cluster on an HPE WDO system. Follow the distribution vendors' instructions for more information on the installation of Hadoop. The reference solution was tested using HDP 2.3 distribution.

For details on setting up SAP HANA Vora components on HPE WDO cluster, refer to Appendix D.

For details on configuring SAP HANA connectivity to SAP HANA Vora on the HPE WDO cluster and configuring SAP HANA Data Lifecycle Manager, refer to Appendix F.

Setting up SAP HANA Vora

Table 1 shows SAP HANA Vora 1.2 components that need to be installed on the HPE EPA cluster.

Table 1. SAP HANA Vora 1.2 components

Software	Description
SAP HANA Vora Base	Base libraries and binaries for SAP HANA Vora
SAP HANA Vora Catalog	Distributed metadata store for SAP HANA Vora
SAP HANA Vora Discovery Service	Manages SAP HANA Vora Service registration
SAP HANA Vora Distributed Log Service	Provides persistence for SAP HANA Vora catalog
SAP HANA Vora Thriftserver	JDBC connectivity for Tools Server, Lumira, etc.
SAP HANA Vora Tools	Web-based UI with SQL editor and OLAP modeler
SAP HANA Vora V2Server	SAP HANA Vora engine

The SAP HANA Vora components need to be installed on appropriate nodes as detailed in the tables below.

HPE recommends the following placement of various Hadoop, Spark and SAP HANA Vora components using the management, compute and storage nodes of the HPE WDO system.

Table 2. Management node base software components

Software	Description
Red Hat® Enterprise Linux® (RHEL) 7.2	Recommended operating system
Oracle Java Development Kit	JDK 1.7.0
Ambari Server, Agent, Management Services, Oozie	Management server and management components for cluster administration
History Server	Maintain information about completed MapReduce jobs
Hive (MetaStore, WebHCat, HiveServer2)	Hive service
Quorum Journal Node	Quorum based Journaling for NameNode High Availability
ZooKeeper	ZooKeeper Service
Apache Spark Client	Apache Spark runtime to initiate Spark YARN client jobs
SAP HANA Vora Base	Base libraries and binaries for SAP HANA Vora
SAP HANA Vora Discovery Server	Manages SAP HANA Vora Service registration
SAP HANA Vora Distributed Log Server	Provides persistence for SAP HANA Vora catalog
SAP HANA Vora Catalog Server	Distributed metadata store for SAP HANA Vora
SAP HANA Vora Thriftserver	JDBC connectivity for Tools Server, Lumira, etc.

Table 3. Head node 1 base software components

Software	Description
Red Hat Enterprise Linux (RHEL) 7.2	Recommended operating system
Oracle Java Development Kit	JDK 1.7.0
Ambari Agent	Ambari Agent on each node for cluster management
NameNode, Quorum Journal Node	HDFS NameNode Service, Quorum based Journaling for NameNode High Availability
ZooKeeper	ZooKeeper Service
Apache Spark Client	Used to initiate Spark YARN-Client jobs (optional)
Spark Job Server	Job Server for Spark applications
SAP HANA Vora Base	Base libraries and binaries for SAP HANA Vora
SAP HANA Vora Discovery Server	Manages SAP HANA Vora Service registration
SAP HANA Vora Distributed Log Server	Provides persistence for SAP HANA Vora catalog
SAP HANA Vora Tools Server	Web-based UI with SQL editor and OLAP modeler

Table 4. Head node 2 base software components

Software	Description
Red Hat Enterprise Linux (RHEL) 7.2	Recommended operating system
Oracle Java Development Kit	JDK 1.7.0
Ambari Agent	Ambari Agent on each node for cluster management
SNameNode	Secondary HDFS NameNode service (on the second Head node)
ResourceManager	YARN resource management service
ZooKeeper	ZooKeeper Service
Apache Spark Client	Used to initiate Spark YARN-Client jobs (optional)
Spark Job Server	Job Server for Spark applications
SAP HANA Vora Base	Base libraries and binaries for SAP HANA Vora
SAP HANA Vora Discovery Server	Manages SAP HANA Vora Service registration
SAP HANA Vora Distributed Log Server	Provides persistence for SAP HANA Vora catalog

Table 5. Compute node base software components

Software	Description
Red Hat Enterprise Linux (RHEL) 7.2	Recommended operating system
Oracle Java Development Kit	JDK 1.7.0
Ambari Agent	Ambari Agent on each node for helping cluster management
NodeManager	NodeManager process for MR2/YARN
Apache Spark	Spark on YARN
SAP HANA Vora Base	Base libraries and binaries for SAP HANA Vora
SAP HANA Vora Discovery Service Client	Manages SAP HANA Vora Service registration
SAP HANA Vora V2Server	SAP HANA Vora engine

Storage node components

Storage nodes contain the software shown in Table 6.

Table 6. Storage node base software components

Software	Description
RHEL 7.2	Recommended operating system
Oracle Java Development Kit	JDK 1.7.0
Ambari Agent	Ambari Agent on each node for helping cluster management
DataNode	DataNode process for HDFS

Configuration changes

Some HDP / Apache Hadoop HDFS/YARN configuration changes are recommended for optimal performance with the HPE EPA/WDO system. Make the changes described here using Ambari.

HDP – HDFS configuration optimizations

Make the following changes to the HDFS configuration.

- Increase the `dfs.blocksize` value to reduce NameNode memory consumption:
`dfs.blocksize 256 or 512` depending on workloads
- Increase the `dfs.namenode.handler.count` value to better manage multiple HDFS operations from multiple clients:
`dfs.namenode.handler.count 180`
- Increase the Java heap size of the NameNode to provide more memory for NameNode metadata:
Java Heap Size of NameNode in Bytes 4096MiB
Java Heap Size of Secondary NameNode in Bytes 4096MiB
- Increase the following timeout value:
`dfs.client.block.write.locateFollowingBlock.retries 30`

HDP – YARN configurations

- Set the amount of physical memory that can be allocated for YARN containers. Allocate about 50% (128 GB) for YARN and the remaining for SAP HANA Vora to allow the HPE EPA cluster to be shared equally by SAP HANA Vora and other Hadoop processing frameworks, e.g., Spark, Hive, etc. If your workload uses SAP HANA Vora predominantly and the SAP HANA Vora in-memory tables need more memory, then the setting should be adjusted downward accordingly.
`yarn.nodemanager.resource.memory-mb 131072MB`
- Set the amount of vcores available on a compute node that can be allocated for containers:
`yarn.nodemanager.resource.cpu-vcores 24`

HDP – Capacity Scheduler

- The `node-locality-delay` specifies how many scheduling intervals to let pass attempting to find a node local slot to run on prior to searching for a rack local slot. This setting is very important for small jobs that do not have a large number of maps or reduces as it will better utilize the compute nodes. We highly recommend this value be set to 1.
`yarn.scheduler.capacity.node-locality-delay 1`

SAP HANA Vora sizing guidelines

Several factors need to be considered when sizing SAP HANA Vora, with the critical factors being data volume and workload characteristics. Data volume impacts the number of compute nodes required, along with storage capacity (both in memory as well as on disk) for the number and size of tables managed in SAP HANA Vora.

Workload characteristics, which are mostly interactive queries, and the expected response times for query execution will dictate the compute requirements.

As SAP HANA Vora is typically deployed to co-exist with Spark, HPE recommends to assign 50% to 80% of the compute node's memory to be used for SAP HANA Vora, although maximizing in-memory access to data, using compute and memory optimized building blocks in the HPE WDO system to manage large data sets completely in memory, provides performance benefits while driving down total cost of ownership.

Refer to SAP's sizing guide for SAP HANA Vora for more detailed guidance on hardware sizing:
http://help.sap.com/Download/Multimedia/hana_Vora/SAP_HANA_Vora_Sizing_Guide_en.pdf

SAP HANA Vora query tuning guidelines

In order to improve the performance of multi-table operations (in particular, JOIN queries), data residing in the SAP HANA Vora engine instances can be partitioned based on one or more columns. You can create a partitioning function using the CREATE PARTITION FUNCTION command. You can create a table that is partitioned by a defined function applied to a table column by using the standard CREATE TABLE statement together with the PARTITIONED BY clause.

When a partitioning function is specified for a column it results in a 1:1 mapping between the underlying SAP HANA Vora engine instances and the column values (the same column value is mapped to the same instance). If the same partitioning function is used for columns from a few different tables, the records from the tables will be co-located, that is, records with the same value for the column for which the partitioning function was specified will be on the same SAP HANA Vora instance. In this case, if a JOIN query involves the columns for which the partitioning function was specified, the whole query will be pushed down to the SAP HANA Vora engine instances and executed there.

There is no difference in the query syntax for co-located and non-co-located tables. However, if a query involves tables where some are co-located and some are not, it should be written in a form that ensures that the operations on co-located tables are executed first. Otherwise, the performance gain from the co-location may not be noticeable.

Depending on the size of the data and its character, the performance benefits from data co-location may not always be noticeable. In general, partitioning should be avoided for small tables (that is, tables which, when loaded, do not encompass all SAP HANA Vora instances in the landscape). As an alternative to boost performance, a replication scheme to replicate small tables is recommended. In general, the number of partitions should be equal to or greater than the number of SAP HANA Vora instances.

Capacity and sizing

Expanding the base configuration of HPE WDO system

As needed, compute and/or storage building blocks can be incrementally added to a base HPE WDO system configuration. The HPE WDO system Reference Architecture, shown in Figure 10, with three HPE Apollo 2000 chassis and four Apollo 4200 Gen9 storage nodes can expand, depending on the requirements, to a variety of combinations of compute and storage nodes in a single rack.

There are a number of options for configuring a single-rack HPE WDO system, 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).

Refer to HPE Reference Architectures for the HPE Elastic Platform for Big Data Analytics (EPA), for detailed sizing and deployment guidance for deploying Hadoop and Spark workloads on your preferred Hadoop distribution.

http://h17007.www1.hp.com/us/en/enterprise/converged-infrastructure/info-library/index.aspx?workload=big_data&type=9

Refer to SAP's sizing guide for SAP HANA Vora for more detailed guidance on hardware sizing:

http://help.sap.com/Download/Multimedia/hana_Vora/SAP_HANA_Vora_Sizing_Guide_en.pdf

Multi-rack configuration

The single-rack HPE WDO system is designed to perform well as a standalone solution but also form the basis of a much larger multi-rack solution, as shown in Figure 23. 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.

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.

Multi-rack extension of single-rack reference architecture

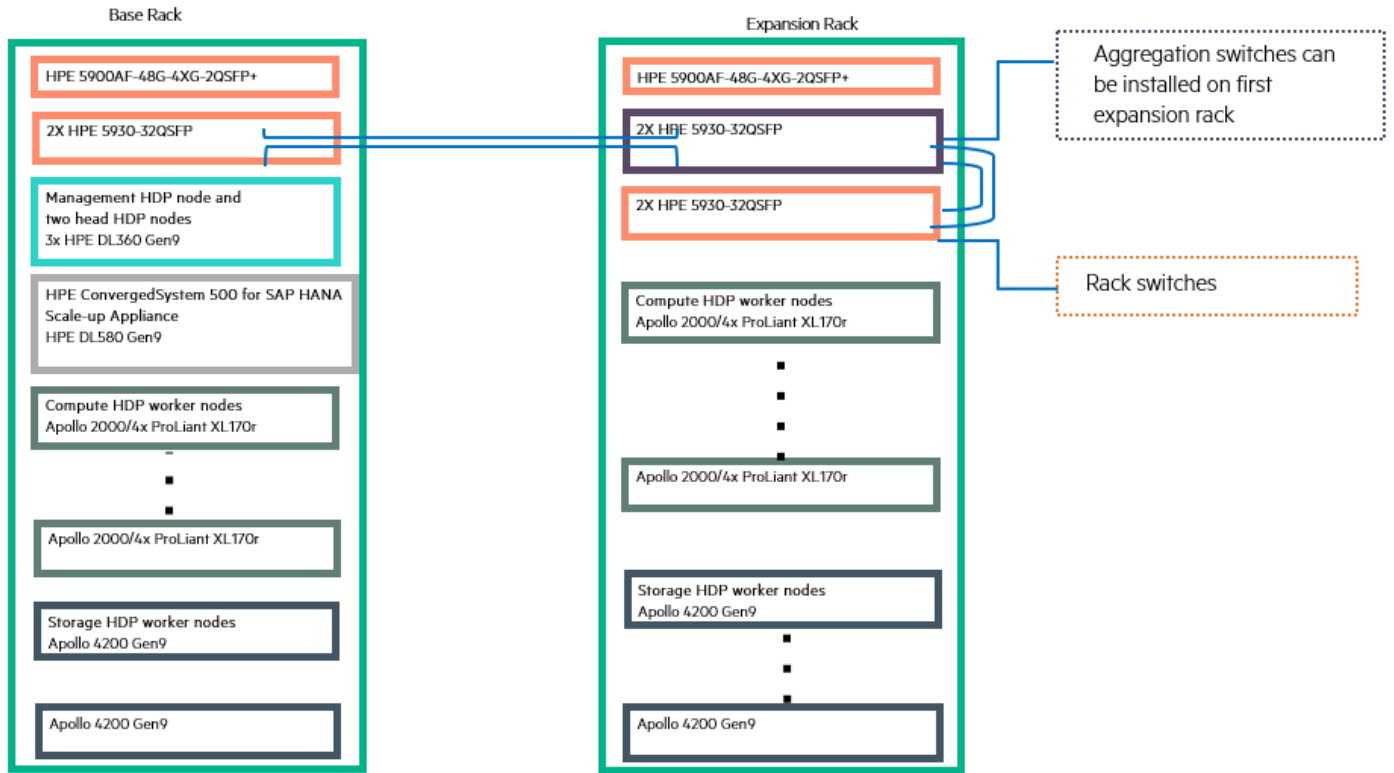


Figure 23. Multi-rack HPE WDO system, extending the capabilities of a single rack

Note

While much of the architecture for the multi-rack solution was borrowed from the single-rack design, the architecture suggested here for multi-rack solutions is based on previous iterations of Hadoop testing on HPE ProLiant servers rather than HPE Apollo servers. It is provided here as a general guideline for designing multi-rack Hadoop clusters.

Extending networking in a multi-rack configuration

For performance and redundancy, two HPE FlexFabric 5930-32QSFP+ ToR switches are specified per expansion rack. The HPE FlexFabric 5930-32QSFP+ switch includes up to six 40 GbE uplinks that can be used to connect these switches to the desired network via a pair of HPE FlexFabric 5930-32QSFP+ aggregation switches.

Guidelines for calculating storage needs

Hadoop cluster storage sizing requires careful planning, including the identification of current and future storage and compute needs. The following are general considerations for data inventory:

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

It makes sense to identify storage requirements for the short-, medium-, and long-term. 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. 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.

Analysis and recommendations

Many organizations are looking to standardize on SAP HANA as a strategic data platform, and integrate Big Data platforms like Hadoop with SAP HANA to enable new analytics capabilities and lower total cost of ownership. The solution described in this Reference Architecture provides deployment and configuration guidance to realize that value using SAP HANA Vora and the HPE Elastic Platform for Big Data Analytics on HDP.

This solution provides enormous flexibility in choice of data processing and storage tier that is based on the value of data, enables a common analytics platform for both data scientists and business analysts, while abstracting the complexity of the Hadoop ecosystem. As a result, Big Data becomes an asset that is managed intelligently to enable new business solutions.

Figure 24 below provides a summary of this solution’s capabilities and its application to appropriate use cases, with cost, volume and performance being key differentiators.

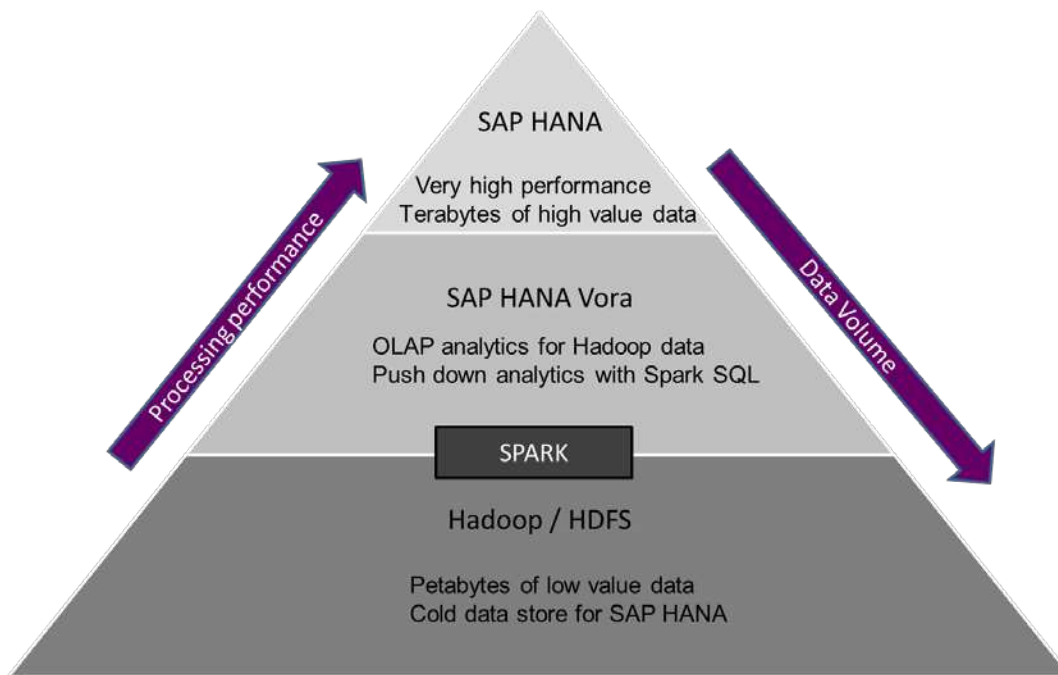


Figure 24. Data management with SAP HANA Vora and Hadoop based on cost, volume and speed

Benefits of HPE Elastic Platform for Big Data Analytics (EPA) using SAP HANA Vora

The HPE Elastic Platform for Big Data Analytics drives down total cost of ownership by enabling consolidation of data, workloads and clusters in a multi-tenant platform. The HPE WDO system provides the flexibility to scale compute and storage independently based on your workload and capacity requirements. Accelerators maximize outcomes, such as accelerating in-memory workloads like Spark and SAP HANA Vora with workload optimized building blocks of HPE Apollo 2000.

Other general benefits from the platform include:

Flexibility to scale

Scale compute and storage independently to build a hyper-scale analytics platform.

The HPE WDO system provides you the flexibility to add modular building blocks that are optimized for each workload and lets you access the same pool of data for batch, interactive or real-time analytics.

Cluster consolidation

Multiple Big Data environments can directly access a shared pool of data.

As Hadoop adoption expands in the enterprise, it is typical to have 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 WDO system allows you to consolidate your data, workloads and clusters on a shared centrally managed platform.

Maximum elasticity

Rapidly provision compute without having to re-partition data.

The HPE WDO system lets you build and scale your infrastructure in alignment with the different phases of your analytics implementation. For example, pilots use a smaller number of storage density optimized servers for data staging and MapReduce workloads, and as you add interactive workloads like SQL on Hadoop or scale your batch workloads, you can add high-latency compute nodes without having to add storage along with compute. A low-latency tier of density optimized servers with SSDs or NVMe Flash can be added to accelerate time series analysis of large datasets in real time.

Breakthrough economics

Lower total cost of ownership through density and workload optimized components.

The HPE WDO system architecture offers 1.6x the compute, 1.5x the storage, and 2.5x memory densities compared to conventional 2U 2-socket servers, and the benefits extend to the hot and cold tiers where you have the flexibility to scale compute or storage based on data access.

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

Implementing a proof-of-concept

As a matter of best practice for all deployments, HPE recommends implementing a proof-of-concept using a test environment that matches as closely as possible the planned production environment. In this way, appropriate performance and scalability characterizations can be obtained. For help with a proof-of-concept, contact an HPE Services representative (hpe.com/us/en/services/consulting.html) or your HPE partner.

Summary

This document provided guidance for deploying SAP HANA Vora using the HPE Workload and Density Optimized (WDO) system. The HPE WDO system uses HPE Apollo 2000 density optimized server building blocks for compute, and HPE Apollo 4200 Gen9 for storage. In addition to simplifying the procurement process, this paper also provided guidelines for integrating SAP HANA Vora with Spark, Hadoop and SAP HANA once the system has been deployed.

Using a predictive maintenance use case, we described the implementation of two key features of SAP HANA Vora:

- Interactive analytics across both SAP HANA and Hadoop data using SAP HANA Vora and Spark
- Manage data movement between SAP HANA and the Hortonworks Data Platform powered by Apache Hadoop

The HPE Apollo 2000 and HPE Apollo 4200 components of the HPE WDO system, help organizations realize breakthrough performance for SAP HANA Vora workloads and enable a flexible and cost-effective approach to consolidating Big Data workloads on the HPE Elastic Platform for Big Data Analytics (EPA).

Appendix A: Bill of materials

The bills of materials (BOMs) provided here are based on a tested configuration for a single-rack HPE EPA solution featuring the following key components:

- One management node
- Two head nodes
- Twelve compute nodes

- Four storage nodes
- Two ToR switches
- One HPE iLO switch
- Hortonworks Data Platform

The following BOMs contain electronic license to use (E-LTU) parts. Electronic software license delivery is now available in most countries. HPE recommends purchasing electronic products over physical products (when available) for faster delivery and for the convenience of not tracking and managing confidential paper licenses. For more information, please contact your reseller or an HPE representative.

Red Hat entitlements

With the Red Hat Enterprise Linux unlimited VM entitlement (G3J24AAE, RHEL Vrtl DC 2 Sckt 3yr 24x7 E-LTU), users have the right to run an unlimited number of desktop or server VMs on each entitled Red Enterprise Virtualization Hypervisor node. As each environment is unique and may not run Red Hat Enterprise Linux exclusively, this entitlement is not included in the parts list for the Reference Architecture. Licensing of the virtual machines is the responsibility of the customer.

Note

Part numbers are at time of publication and subject to change. The bill of materials does not include complete support options or other rack and power requirements. If you have questions regarding ordering, please consult with your HPE Reseller or HPE Sales Representative for more details. hpe.com/us/en/services/consulting.html

Management node and head nodes

Important

Table A-1 provides a BOM for one HPE ProLiant DL360 Gen9 server. The tested solution featured one management server and two head nodes, requiring a total of three HPE ProLiant DL360 Gen9 servers.

Table A-1. BOM for a single HPE ProLiant DL360 Gen9 server

Qty	Part number	Description
1	755258-B21	HPE DL360 Gen9 8SFF CTO Server
1	755386-L21	HPE DL360 Gen9 E5-2640v3 FIO Kit
1	755386-B21	HPE DL360 Gen9 E5-2640v3 Kit
8	726719-B21	HPE 16GB 2Rx4 PC4-2133P-R Kit
8	652589-B21	HPE 900GB 6G SAS 10K 2.5in SC ENT HDD
1	749974-B21	HPE Smart Array P440ar/2G FIO Controller
1	779799-B21	HPE Ethernet 10Gb 2P 546FLR-SFP+ Adptr
1	734807-B21	HPE 1U SFF Easy Install Rail Kit
2	720478-B21	HPE 500W FS Plat Ht Plg Pwr Supply Kit
1	764646-B21	HPE DL360 Gen9 Serial Cable
1	764636-B21	HPE DL360 Gen9 SFF Sys Insght Dsply Kit
2	AF595A	HPE 3.0M, Blue, CAT6 STP, Cable Data

Compute nodes

Important

Table A-2 provides a BOM for one HPE Apollo r2600 24SFF chassis with four HPE ProLiant XL170r servers. The reference solution features three HPE Apollo r2600 chassis with a total of 12 HPE ProLiant XL170r servers. The tested solution used 12 HPE ProLiant XL170r servers.

Table A-2. BOM for a single HPE Apollo r2600 24SFF chassis with 4 HPE ProLiant XL170r servers

Qty	Part number	Description
1	798153-B21	HPE Apollo r2600 24SFF CTO Chassis
2	800059-B21	HPE Apollo 2000 FAN-module Kit
4	798155-B21	HPE ProLiant XL170r Gen9 CTO Svr
4	793028-L21	HPE XL1x0r Gen9 E5-2680v3 FIO Kit
4	793028-B21	HPE XL1x0r Gen9 E5-2680v3 Kit
64	726719-B21	HPE 16GB 2Rx4 PC4-2133P-R Kit
4	779799-B21	HPE Ethernet 10Gb 2P 546FLR-SFP+ Adptr
8	832414-B21	HPE 480GB 6Gb SATA 2.5in MU-2 SC SSD
4	798178-B21	HPE XL1x0r Gen9 LP PCIe16 L Riser Kit
4	798180-B21	HPE XL170r FLOM x8 R Riser Kit
4	798192-B21	HPE XL170r/190r Dedicated NIC IM Board Kit
4	800060-B21	HPE XL170r Mini-SAS B140 Cbl Kit
2	720620-B21	HPE 1400W FS Plat Pl Ht Plg Pwr Supply Kit
1	740713-B21	HPE t2500 Strap Shipping Bracket
1	611428-B21	HPE DL2000 Hardware Rail Kit
8	AF595A	HPE 3.0M, Blue, CAT6 STP, Cable Data
2	JG330A	HPE X240 QSFP+ 4x10G SFP+ 3m DAC Cable

Storage nodes

Note

Table A-3 provides a BOM for one HPE Apollo 4200 Gen9 server. The reference solution features three to five HPE Apollo 4200 Gen9 servers. The tested solution used four HPE Apollo 4200 Gen9 servers.

Table A-3. BOM for a single HPE Apollo 4200 Gen9 server

Qty	Part number	Description
1	808027-B21	HPE Apollo 4200 Gen9 24LFF CTO Svr
1	803311-L21	HPE Apollo 4200 Gen9 E5-2660v3 FIO Kit
1	803311-B21	HPE Apollo 4200 Gen9 E5-2660v3 Kit
8	726719-B21	HPE 16GB 2Rx4 PC4-2133P-R Kit
1	806563-B21	HPE Apollo 4200 Gen9 LFF Rear HDD Cage Kit
28	797265-B21	HPE 4TB 6G SATA 7.2k 3.5in MDL LP HDD
1	813546-B21	HPE SAS Controller Mode for Rear Storage
1	764285-B21	HPE IB FDR/EN 40Gb 2P 544_FLR-QSFP Adapter
2	720479-B21	HPE 800W FS Plat Hot Plug Power Supply Kit
1	806565-B21	HPE Apollo 4200 Gen9 IM Card Kit
1	788028-B21	HPE 120GB RI Solid State M.2 Kit
1	806562-B21	HPE Apollo 4200 Gen9 Redundant Fan Kit
2	JG327A	HPE X240 40G QSFP+ QSFP+ 3m DAC Cable
2	AF595A	HPE 3.0M, Blue, CAT6 STP, Cable Data

Networking

Table A-4 provides a BOM for two ToR switches and one iLO switch, as featured in the tested configuration.

Table A-4. BOM for two HPE FlexFabric 5930 switches (ToR) and one HPE FlexFabric 5900 switch (HPE iLO)

Qty	Part number	Description
2	JG726A	HPE FF 5930-32QSFP+ Switch
4	JG553A	HPE X712 Bck(pwr)-Frt(prt) HV Fan Tray
4	JC680A	HPE A58x0AF 650W AC Power Supply
1	JG510A	HPE 5900AF-48G-4XG-2QSFP+ Switch
2	JC680A	HPE A58x0AF 650W AC Power Supply
2	JC682A	HPE 58x0AF Bck(pwr)-Frt(ports) Fan Tray
2	JG329A	HPE X240 QSFP+ 4x10G SFP+ 1m DAC Cable
4	JG326A	HPE X240 40G QSFP+ QSFP+ 1m DAC Cable

Other hardware

Table A-5. BOM for a single rack with four PDUs

Qty	Part number	Description
1	BW904A	HPE 42U 600x1075mm Enterprise Shock Rack
1	BW946A	HPE 42U Location Discovery Kit
1	BW930A	HPE Air Flow Optimization Kit
1	TK817A	HPE CS Rack Side Panel 1075mm Kit
1	TK816A	HPE CS Rack Light Kit
1	TK815A	HPE CS Rack Door Branding Kit
1	BW891A	HPE Rack Grounding Kit
4	AF520A	HPE Intelligent Mod PDU 24a Na/Jpn Core
6	AF547A	HPE 5xC13 Intlght PDU Ext Bars G2 Kit
2	C7536A	HPE Ethernet 14ft CAT5e RJ45 M/M Cable

Recommended service components

Table A-6 provides a BOM for three recommended service components for Factory Express Build, TS Consulting and HPE WDO system.

Table A-6. Recommended service components

Qty	Part number	Description
--	HA454A1	HPE Factory Express Level 4 Service (recommended)
1	H8E04A1	HPE Hadoop Custom Consulting Service (recommended)
1	P6L57A	HPE Big Data Reference Architecture

Software options

Important

Quantities listed in Table A-7 may vary. Quantities below are based on a rack with three management/head nodes, twelve compute nodes and four storage nodes.

Table A-7. BOM for software options

Qty	Part number	Description
18	E6U59ABE	HPE iLO Adv incl 1yr TS U E-LTU
18	QL803BAE	HPE Insight CMU 1yr 24x7 Flex E-LTU
18	G3J29AAE	RHEL Svr 2 Sckt/2 Gst 1yr 9x5 E-LTU
3	C9A82AAE	HPE Secure Encryption per Server Entitlement

Hortonworks software

Important

Table A-8 provides the BOM for the Hortonworks license and support. One Hortonworks license covers 1 HPE Apollo 2000 chassis of compute nodes.

For the HPE Moonshot compute option, one Hortonworks license covers up to 16 HPE ProLiant m710p server cartridges. Therefore, three of these licenses will cover a fully populated HPE Moonshot 1500 chassis.

While HPE is a certified reseller of Hortonworks software subscriptions, all application support (level-one through level-three) is provided by Hortonworks.

Table A-8. BOM for Hortonworks software

Qty	Part number	Description
--	F5Z52A	Hortonworks Data Platform Enterprise 4 Nodes or 50TB Raw Storage 1 year 24x7 Support LTU

Appendix B: Alternate HPE EPA components for compute and storage

This appendix provides BOMs for alternate processors, memory, and disk drives for the HPE ProLiant XL170r servers used as compute nodes and for the HPE Apollo 4200 servers used as storage nodes.

Alternate compute node components

Table B-1. Alternate processors – HPE Apollo 2000 – HPE ProLiant XL170r

Qty	Part number	Description
1	793028-L21	HPE XL1x0r Gen9 E5-2680v3 FIO Kit (12C, 2.5 GHz)
1	793028-B21	HPE XL1x0r Gen9 E5-2680v3 Kit
1	793024-L21	HPE XL1x0r Gen9 E5-2660v3 FIO Kit (10C, 2.6 GHz)
1	793024-B21	HPE XL1x0r Gen9 E5-2660v3 Kit
1	793020-L21	HPE XL1x0r Gen9 E5-2640v3 FIO Kit (8C, 2.6 GHz)
1	793020-B21	HPE XL1x0r Gen9 E5-2640v3 Kit

Table B-2. Alternate memory – HPE Apollo 2000 – HPE ProLiant XL170r

Qty per node	Part number	Description
8/16	728629-B21	HPE 32GB 2Rx4 PC4-2133P-R Kit
8/16	726719-B21	HPE 16GB 2Rx4 PC4-2133P-R Kit
16	759934-B21	HPE 8GB 2Rx8 PC4-2133P-R Kit

Table B-3. Alternate disk drives – HPE Apollo 2000 – HPE ProLiant XL170r

Qty per node	Part number	Description
2/4/6	804625-B21	HPE 800GB 6G SATA MU-2 SFF SC SSD
2/4/6	817011-B21	HPE 1.92TB 6G SATA MU-3 SFF SC SSD
2/4/6	757339-B21	HPE 1.6TB 6G SATA VE 2.5in SC EV SSD
1	798190-B21	HPE XL1x0r Gen9 NGFF Riser w/ 2x64G Drive (Optional for Operating System)

Alternate storage node components

Table B-4. Alternate processors – HPE Apollo 4200

Qty	Part number	Description
1	821791-L21	HPE Apollo 4200 Gen9 E5-2697v3 FIO Kit (14C, 2.6GHz)
1	821791-B21	HPE Apollo 4200 Gen9 E5-2697v3 Kit
1	803314-L21	HPE Apollo 4200 Gen9 E5-2680v3 FIO Kit (12C, 2.5 GHz)
1	803314-B21	HPE Apollo 4200 Gen9 E5-2680v3 Kit
1	803308-L21	HPE Apollo 4200 Gen9 E5-2640v3 FIO Kit (8C, 2.6 GHz)
1	803308-B21	HPE Apollo 4200 Gen9 E5-2640v3 Kit

Table B-5. Alternate memory – HPE Apollo 4200 (Quantity per node specific to customer deployment)

Qty per node	Part number	Description
Deployment-specific	728629-B21	HPE 32GB 2Rx4 PC4-2133P-R Kit
Deployment-specific	726717-B21	HPE 4GB 1Rx8 PC4-2133P-R Kit
Deployment-specific	759934-B21	HPE 8GB 2Rx8 PC4-2133P-R Kit

Appendix C: HPE value added services and support

In order to help you jump-start your Hadoop solution development, HPE offers a range of Big Data services, which are outlined in this appendix.

Factory Express Services

Factory-integration services are available for customers seeking a streamlined deployment experience. With the purchase of Factory Express services, your Hadoop cluster will arrive racked and cabled, with software installed and configured per an agreed-upon custom statement of work, for the easiest deployment possible. You should contact HPE Technical Services for more information and for assistance with a quote.

Technical Services Consulting – Reference Architecture Implementation Service for Hadoop (Hortonworks)

With HPE Reference Architecture Implementation Service for Hadoop, HPE can install, configure, deploy, and test a Hadoop cluster that is based on HPE EPA. Experienced consultants implement all the details of the original Hadoop design: naming, hardware, networking, software, administration, backup, disaster recovery, and operating procedures. Where options exist, or the best choice is not clear, HPE works with you to configure the environment to meet your goals and needs. HPE also conducts an acceptance test to validate that the system is operating to your satisfaction.

Technical Services Consulting – Big Data Services

HPE Big Data Services can help you reshape your IT infrastructure to corral increasing volumes of data – from e-mails, social media, and website downloads – and convert them into beneficial information. These services encompass strategy, design, implementation, protection, and compliance. Delivery is in the following three steps:

1. **Architecture strategy:** HPE defines the functionalities and capabilities needed to align your IT with your Big Data initiatives. Through transformation workshops and roadmap services, you'll learn to capture, consolidate, manage and protect business-aligned information, including structured, semi-structured, and unstructured data.
2. **System infrastructure:** HPE designs and implements a high-performance, integrated platform to support a strategic architecture for Big Data. Choose from design and implementation services, Reference Architecture implementations, and integration services. Your flexible, scalable infrastructure will support Big Data variety, consolidation, analysis, share, and search on HPE platforms.
3. **Data protection:** Ensure the availability, security, and compliance of your Big Data systems. HPE can help you safeguard your data and achieve regulatory compliance and lifecycle protection across your Big Data landscape, while also enhancing your approach to backup and business continuity.

For additional information, visit hpe.com/us/en/services/consulting/big-data.html.

HPE Support options

HPE offers a variety of support levels to meet your needs. More information is provided below.

HPE Support Plus 24

HPE can provide integrated onsite hardware/software support services, available 24x7x365, including access to HPE technical resources, four-hour response onsite hardware support and software updates.

HPE Proactive Care

HPE Proactive Care provides all of the benefits of proactive monitoring and reporting, along with rapid reactive support through HPE's expert reactive support specialists. You can customize your reactive support level by selecting either six-hour call-to-repair or 24x7 with four-hour onsite response.

HPE Proactive Care helps prevent problems, resolve problems faster, and improve productivity. Through analysis, reports, and update recommendations, you are able to identify and address IT problems before they can cause performance issues or outages.

HPE Proactive Care with the HPE Personalized Support Option

Adding the Personalized Support Option for HPE Proactive Care is highly recommended. This option builds on the benefits of HPE Proactive Care Service, providing you an assigned Account Support Manager who knows your environment and can deliver support planning, regular reviews, and technical and operational advice specific to your environment.

HPE Proactive Select

To address your ongoing/changing needs, HPE recommends adding Proactive Select credits to provide tailored support options from a wide menu of services that can help you optimize the capacity, performance, and management of your environment. These credits may also be used for assistance in implementing solution updates. As your needs change over time, you have the flexibility to choose the services best suited to address your current challenges.

HPE Datacenter Care

HPE Datacenter Care provides a more personalized, customized approach for large, complex environments, providing a single solution for reactive, proactive, and multi-vendor support needs. You may also choose the Defective Media Retention (DMR) option.

Other offerings

HPE highly recommends HPE Education Services (customer training and education) and additional Technical Services, as well as in-depth installation or implementation services when needed.

More information

For additional information, visit:

- HPE Education Services: <http://h10076.www1.hp.com/ww/en/training/portfolio/bigdata.html>
- HPE Technology Consulting Services: hpe.com/us/en/services/consulting/big-data.html
- HPE Deployment Services: hpe.com/services/deployment

Appendix D: Installing SAP HANA Vora on HPE EPA cluster

Installing HDP 2.3 and Spark

Follow the standard installation procedure for setting up the HPE EPA cluster for HDP 2.3 before installing the SAP HANA Vora components.

HDP 2.3 comes with Spark version 1.3 by default. While this is production ready and supported, SAP HANA Vora 1.2 is compatible only with Spark release version 1.5.2. Download this version of Spark from Apache Spark website (<https://spark.apache.org>).

- Choose the package type as **Pre-built for Hadoop 2.6** and download the `spark-1.5.2-bin-hadoop2.6.tgz` file.
- Extract the contents of this file to the `/opt/spark` directory on the management node (`r01mgt`) and all the compute nodes (`r01wn01...12`).

- Set the following environment variables in appropriate startup scripts on all the nodes where Spark 1.5.2 is installed.

```
export SPARK_HOME=/opt/spark/spark-1.5.2-bin-hadoop2.6
```

```
export SPARK_CONF_DIR=$SPARK_HOME/conf
```

```
export PATH=$SPARK_HOME/bin:$PATH
```

- In the `/opt/spark/spark-1.5.2-bin-hadoop2.6` directory, copy the `spark-defaults.conf.template` file as `spark-defaults.conf` and set the following properties:

```
spark.master yarn-client
```

```
spark.driver.extraJavaOptions -Dhdp.version=2.3.0.0-2557
```

```
spark.yarn.am.extraJavaOptions -Dhdp.version=2.3.0.0-2557
```

Installing and configuring SAP HANA Vora components

The SAP HANA Vora components need to be installed and started in a way that enables them to operate together correctly. Follow the following steps to ensure that SAP HANA Vora is set up correctly.

Installing prerequisite software

SAP HANA Vora Distributed Log component requires the `libaio` package to be installed on all the nodes where this component is installed. Install the `libaio` package on the management node and on the two head nodes. Run the following command as root user:

```
# yum install libaio
```

Installing SAP HANA Vora Ambari package

The SAP HANA Vora Ambari package can be downloaded from the SAP software download center (<https://support.sap.com/swdc>). You'll need an S-User ID to log in to SAP Service Marketplace.

- Download the `VORA_AM1_02_0-70001227.GZ` zip file from SAP software download center to the management node (r01mgt).
- Extract the zip file to the directory `/var/lib/ambari-server/resources/stacks/HDP/2.3/services`
- Restart Ambari server with the following command:

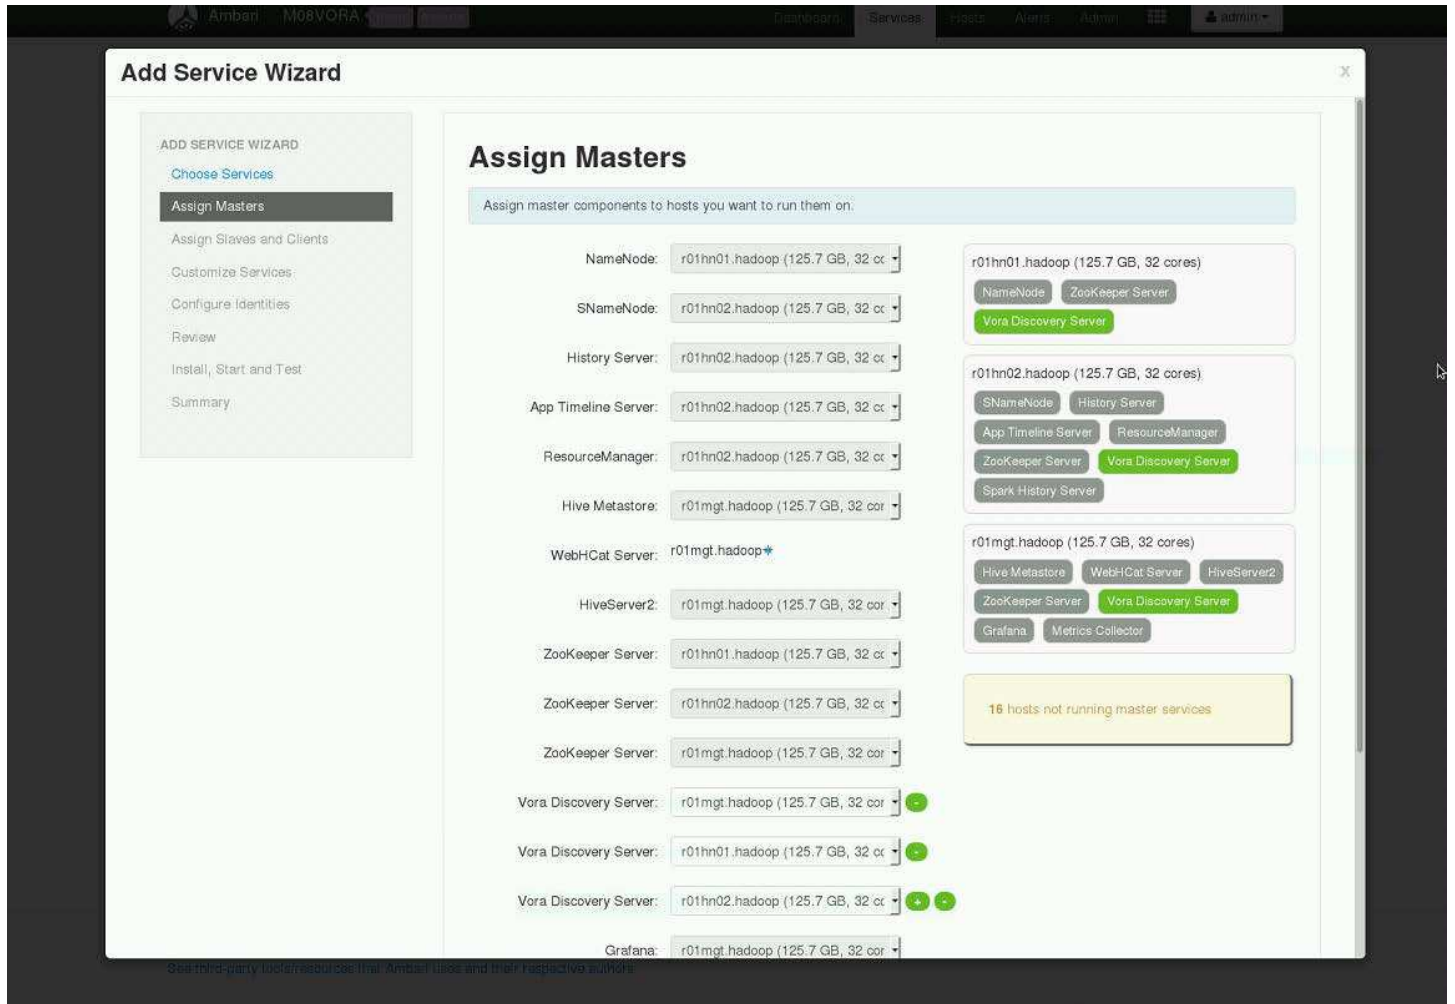
```
# ambari-server restart
```

Installing SAP HANA Vora Base

- From the Ambari Dashboard, click the **Actions** button, then choose **+ Add Service** to invoke the Add Service wizard.
- On the **Choose Services** screen, select the **Vora Base** option and click **Next**.
- On the **Assign Slaves and Clients** screen, choose all nodes and click **Next**.
- There is no customization needed for **Vora Base** service.

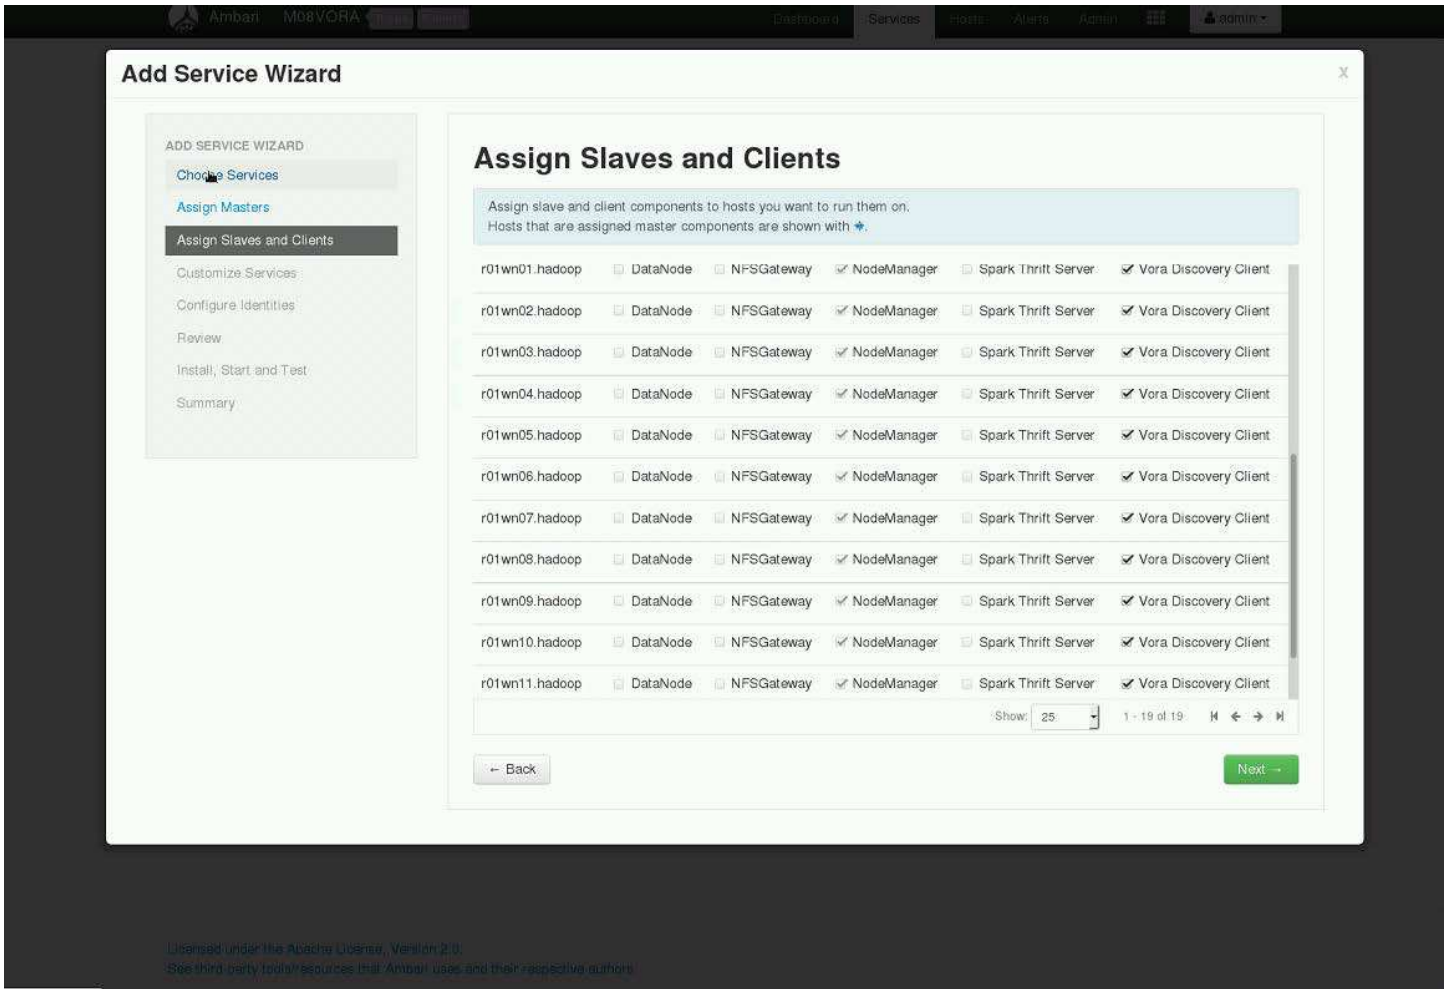
Installing SAP HANA Vora Discovery Service

- From the Ambari Dashboard, click the **Actions** button, then choose **+ Add Service** to invoke the Add Service wizard.
- On the **Choose Services** screen, select the **Vora Discovery** option and click **Next**.
- On the **Assign Masters** screen, choose the management and head nodes as shown in the figure and click **Next**.



Some third-party logos are used here that Ambari does not own and their respective authors.

- On the **Assign Slaves and Clients** screen, choose all the compute nodes (r01wn01...12) as shown in the figure and click **Next**.



- On the **Customize Services** screen, under the **Advanced vora-discovery-config** section, enter the values for the parameters as shown in the figure:

The screenshot shows a configuration interface for 'Advanced vora-discovery-config'. It contains several input fields with the following parameters and values:

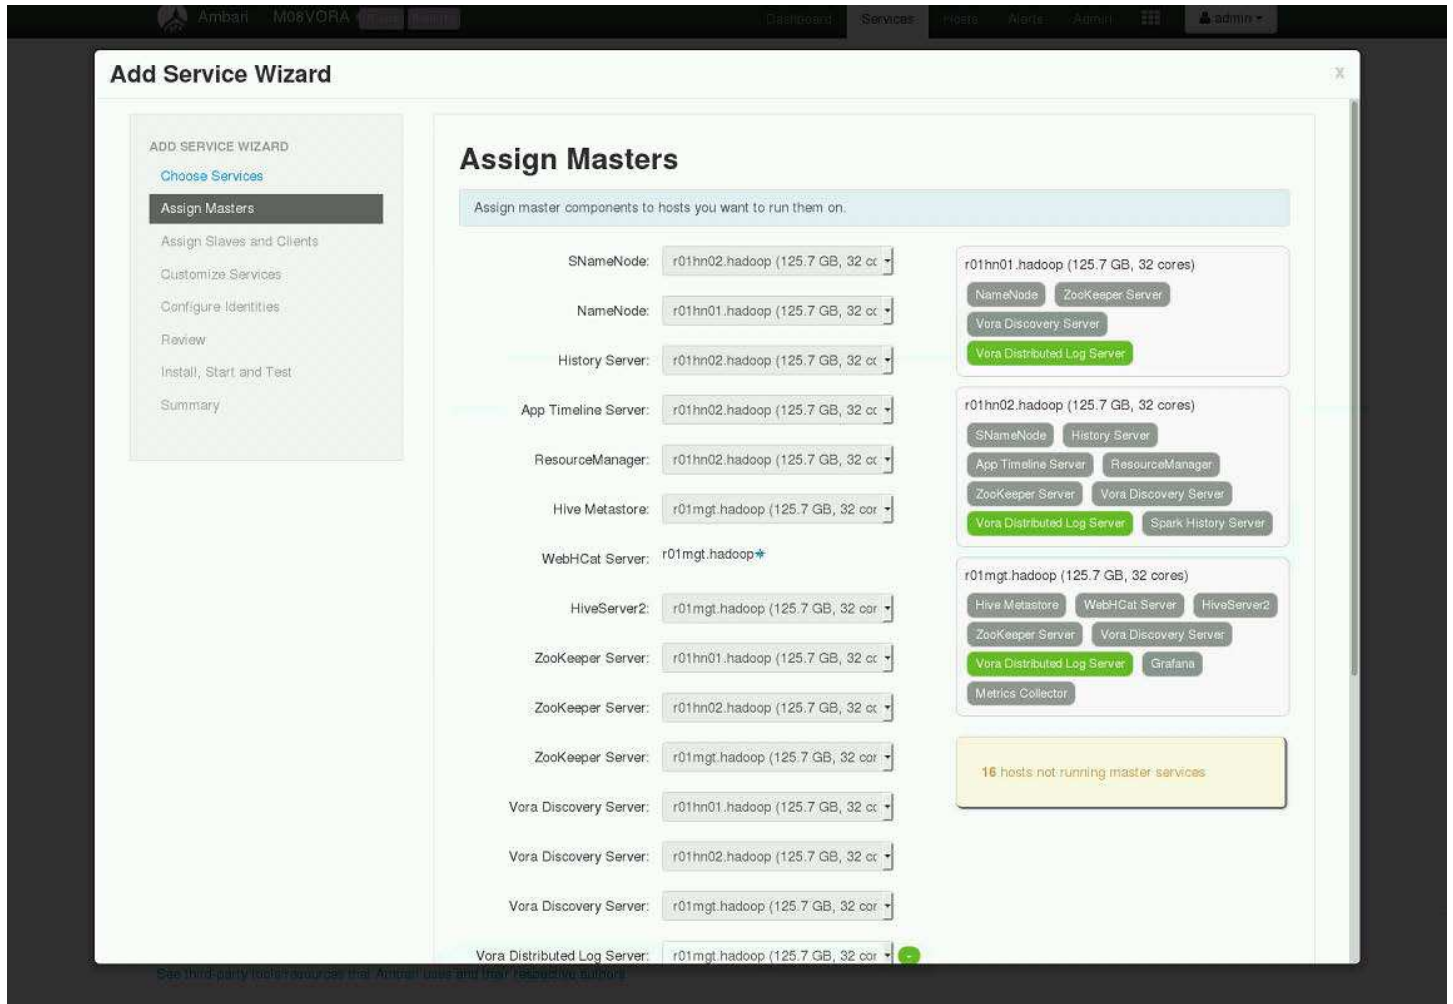
Parameter	Value
vora_discovery_bind_interface	bond0
vora_discovery_bootstrap_host	r01mgt.hadoop
vora_discovery_bootstrap_timeout	30
vora_discovery_data_dir	/var/local/vora-discovery
vora_discovery_log_dir	/var/log/vora-discovery
vora_discovery_log_level	WARNING
vora_discovery_servers	r01hn01.hadoop,r01hn02.hadoop

Below the configuration fields, there is a section for 'Custom vora-discovery-config' which is currently collapsed.

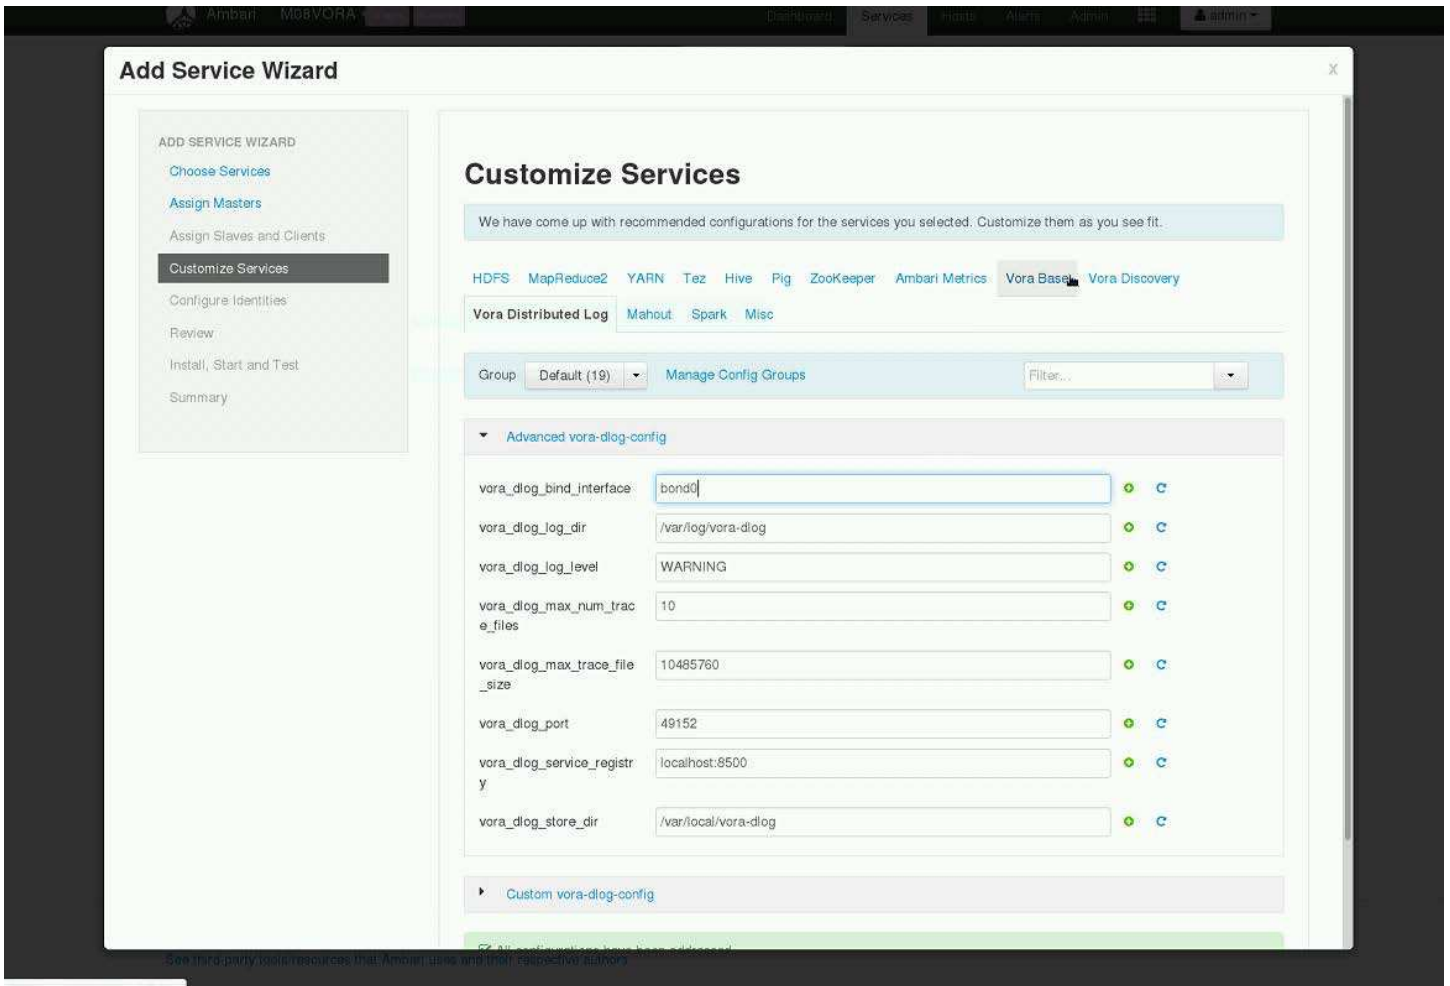
- Finish the installation after reviewing the selections.

Installing SAP HANA Vora Distributed Log Service

- From the Ambari Dashboard, click the **Actions** button, then choose **+ Add Service** to invoke the Add Service wizard.
- On the **Choose Services** screen, select the **Vora Distributed Log** option and click **Next**.
- On the **Assign Masters** screen, choose the management and head nodes as shown in the figure and click **Next**.



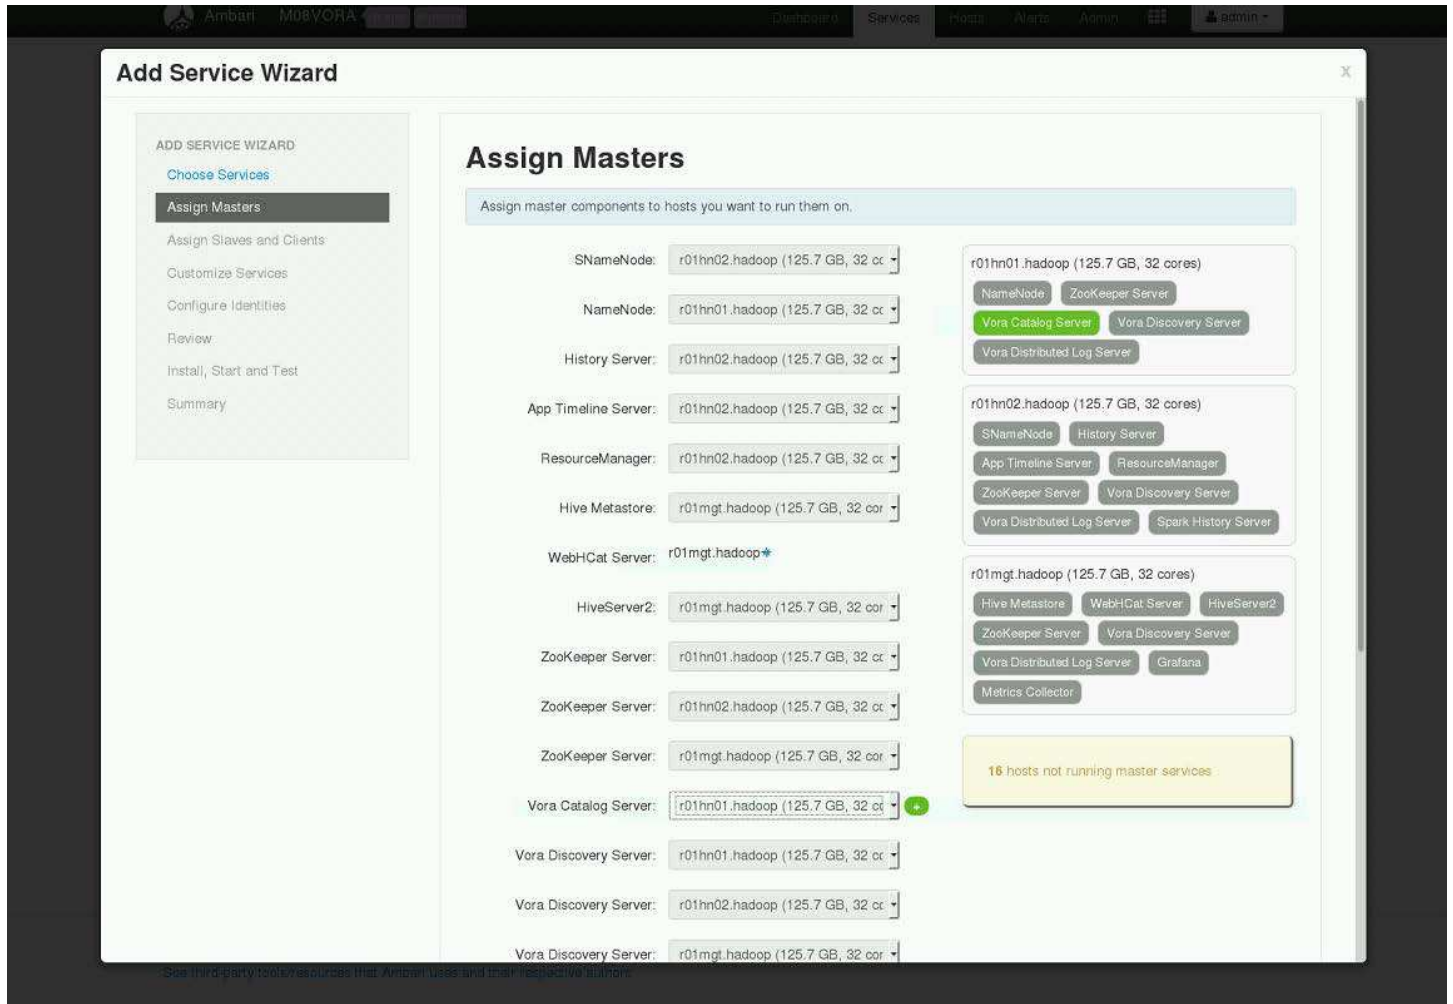
- On the **Customize Services** screen, under the **Advanced vora-dlog-config** section, enter the values for the parameters as shown in the figure:



- Finish the installation after reviewing the selections.

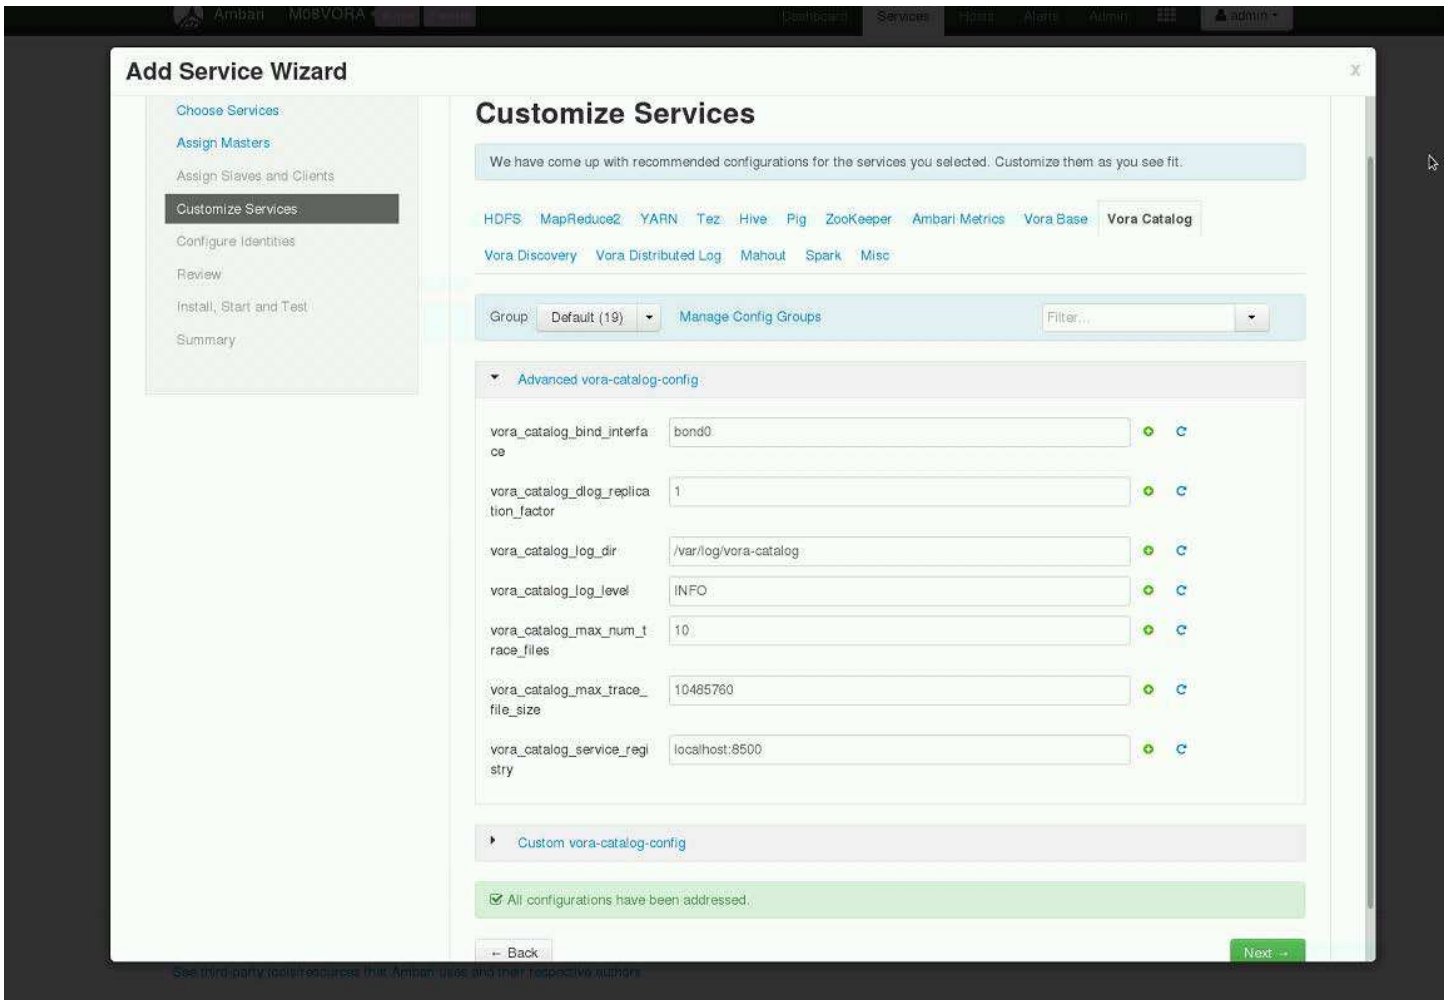
Installing SAP HANA Vora Catalog Service

- From the Ambari Dashboard, click the **Actions** button, then choose **+ Add Service** to invoke the Add Service wizard.
- On the **Choose Services** screen, select the **Vora Catalog** option and click **Next**.
- On the **Assign Masters** screen, choose the head node **r01hn01** as shown in the figure and click **Next**.



See third-party trademarks that Ambari uses and their respective owners.

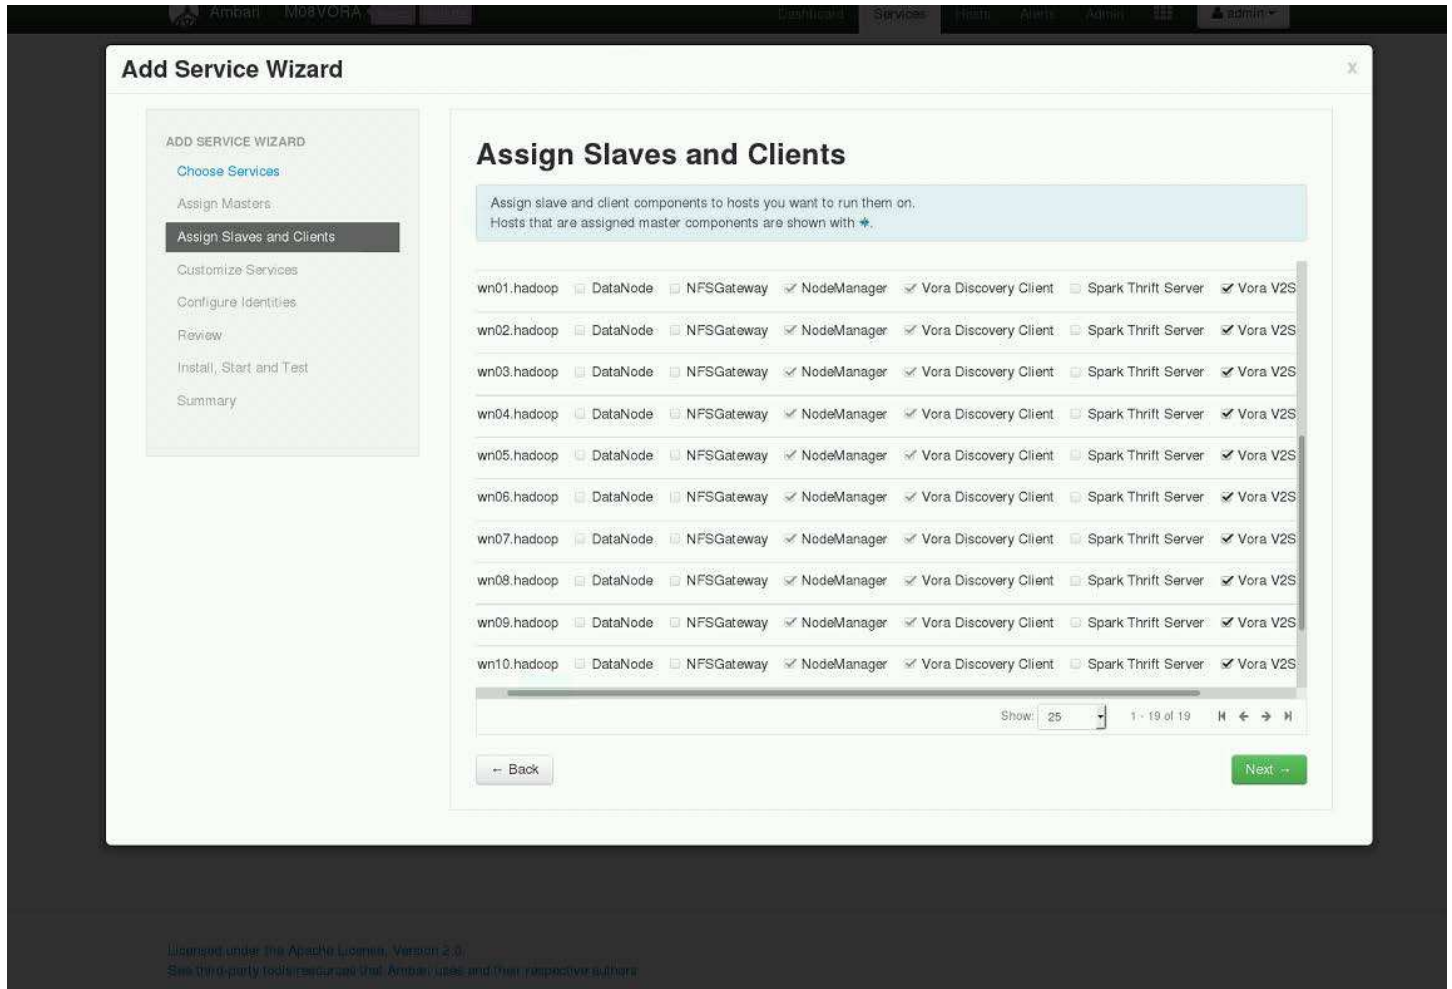
- On the **Customize Services** screen, under the **Advanced vora-catalog-config** section, enter the values for the parameters as shown in the figure:



- Finish the installation after reviewing the selections.

Installing SAP HANA Vora V2Server Service

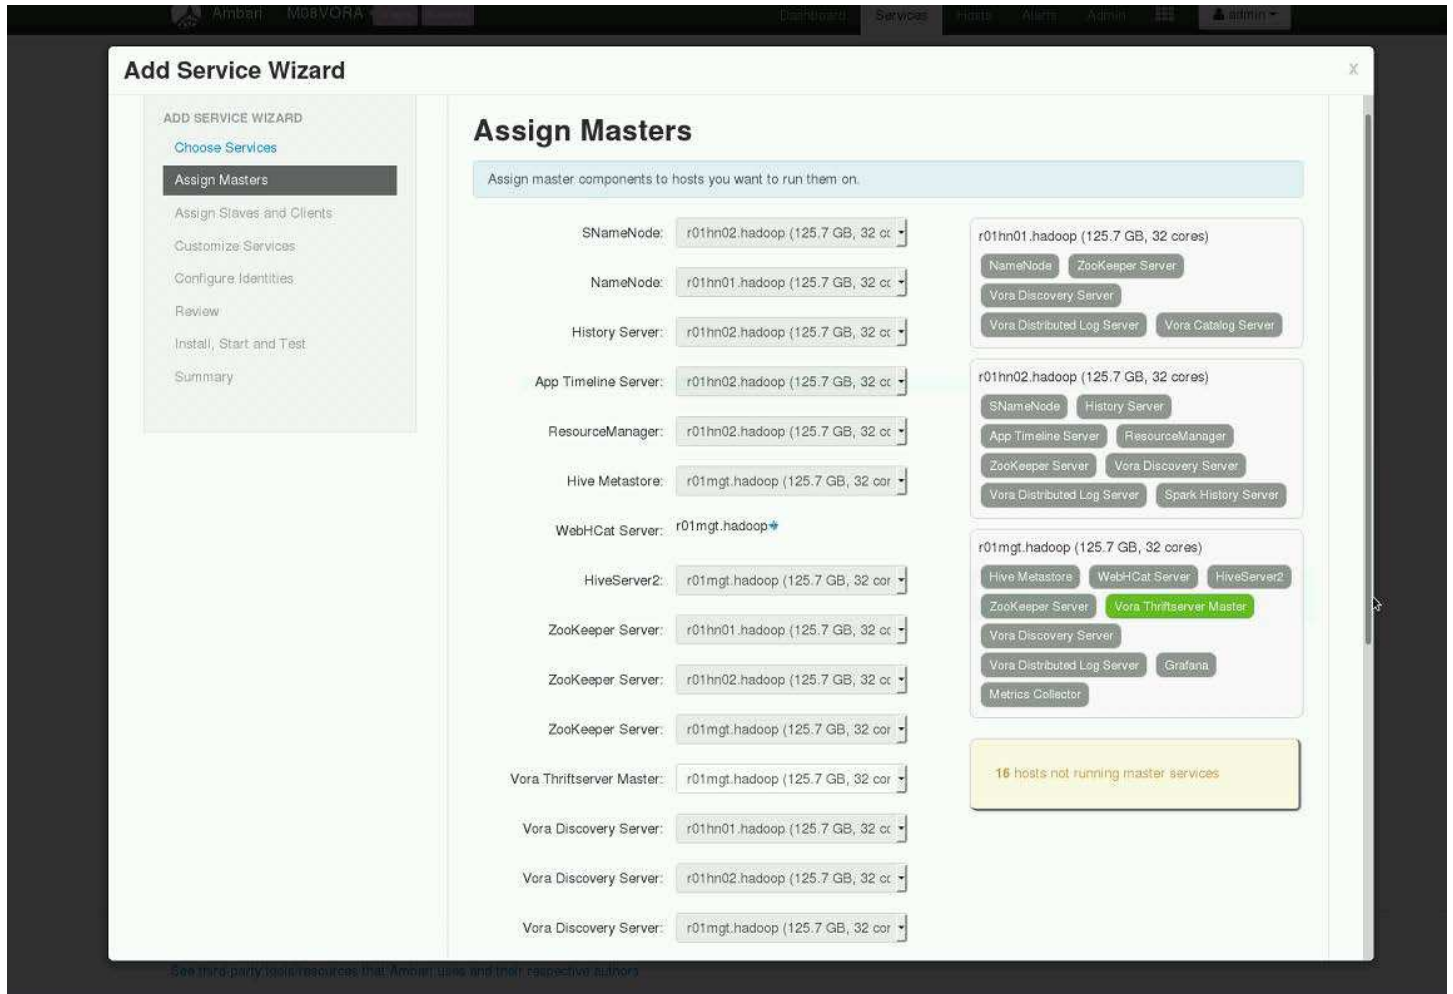
- From the Ambari Dashboard, click the **Actions** button, then choose **+ Add Service** to invoke the Add Service wizard.
- On the **Choose Services** screen, select the **Vora V2Server** option and click **Next**.
- On the **Assign Slaves and Clients** screen, choose all compute nodes (**r01wn01...12**) as shown in the figure and click **Next**.



- On the **Customize Services** screen, choose default values and click **Next**.
- Finish the installation after reviewing the selections.

Installing SAP HANA Vora Thriftserver Service

- From the Ambari Dashboard, click the **Actions** button, then choose **+ Add Service** to invoke the Add Service wizard.
- On the **Choose Services** screen, select the **Vora Thriftserver** option and click **Next**.
- On the **Assign Masters** screen, choose the management node **r01mgt** as shown in the figure and click **Next**.



- On the **Customize Services** screen, under the **Advanced vora-thriftserver-config** section, enter the values for the parameters as shown in the figure:

Advanced vora-thriftserver-config

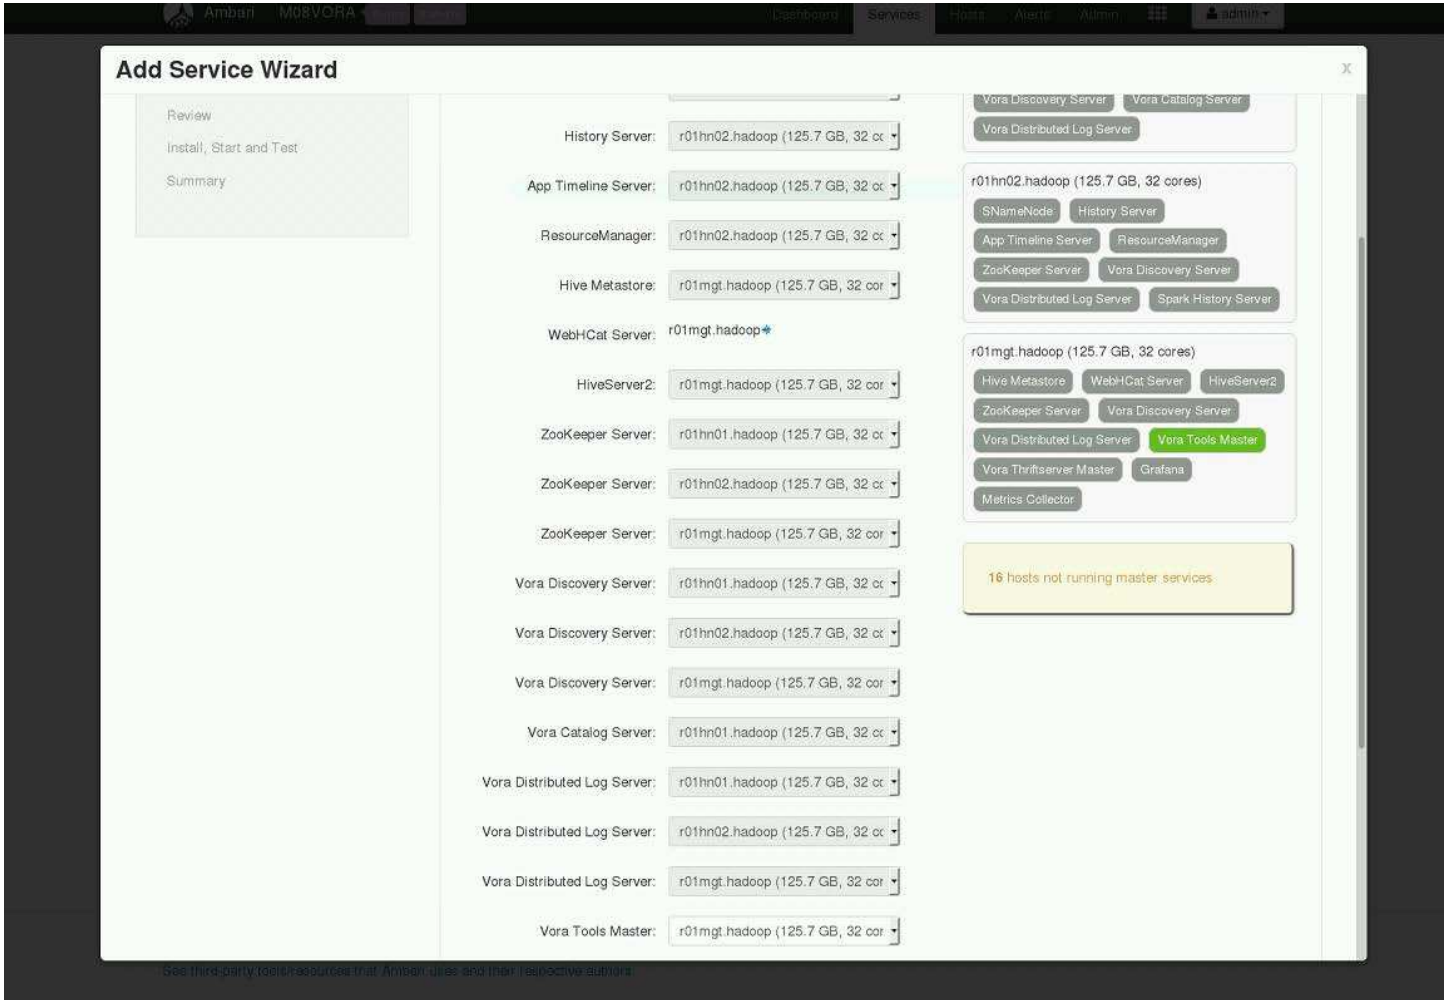
vora_thriftserver_extra_arguments	<input type="text" value="--conf vora.hadoop.distro=hdp"/>	⊗
vora_thriftserver_java_home	<input type="text" value="/usr/java/jdk1.7.0_75/"/>	⊗
vora_thriftserver_log_dir	<input type="text" value="/var/log/vora-thriftserver"/>	⊗
vora_thriftserver_log_level	<input type="text" value="WARNING"/>	⊗
vora_thriftserver_metastore_dir	<input type="text" value="/tmp/vora-thriftserver"/>	⊗
vora_thriftserver_service_registry	<input type="text" value="localhost:8500"/>	⊗
vora_thriftserver_spark_home	<input type="text" value="/opt/spark/spark-1.5.2-bin-hadoop2.6/"/>	⊗

Custom vora-thriftserver-config

- Finish the installation after reviewing the selections.

Installing SAP HANA Vora Tools Service

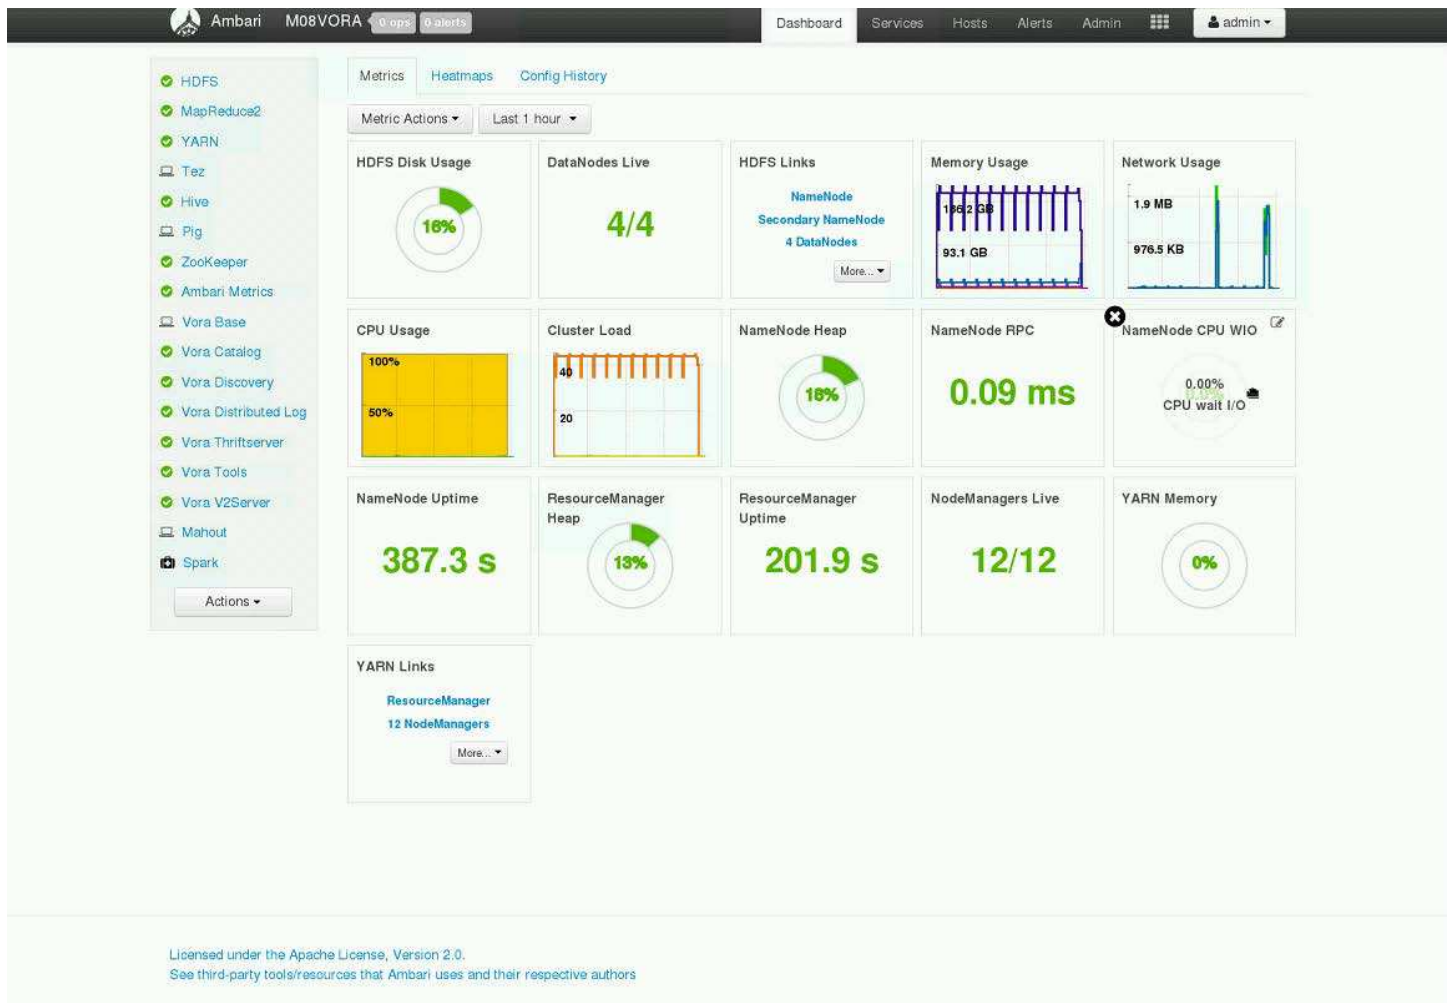
- From the Ambari Dashboard, click the **Actions** button, then choose **+ Add Service** to invoke the Add Service wizard.
- On the **Choose Services** screen, select the **Vora Tools** option and click **Next**.
- On the **Assign Masters** screen, choose the management node **r01mgt** as shown in the figure and click **Next**.



- On the **Customize Services** screen, choose default values and click **Next**.
- Finish the installation after reviewing the selections.

Verify the SAP HANA Vora components installation

After all the above services are installed, go to the Ambari Dashboard and make sure that all the SAP HANA Vora services are installed without errors and are up as shown in the following figure:



Licensed under the Apache License, Version 2.0.
See third-party tools/resources that Ambari uses and their respective authors

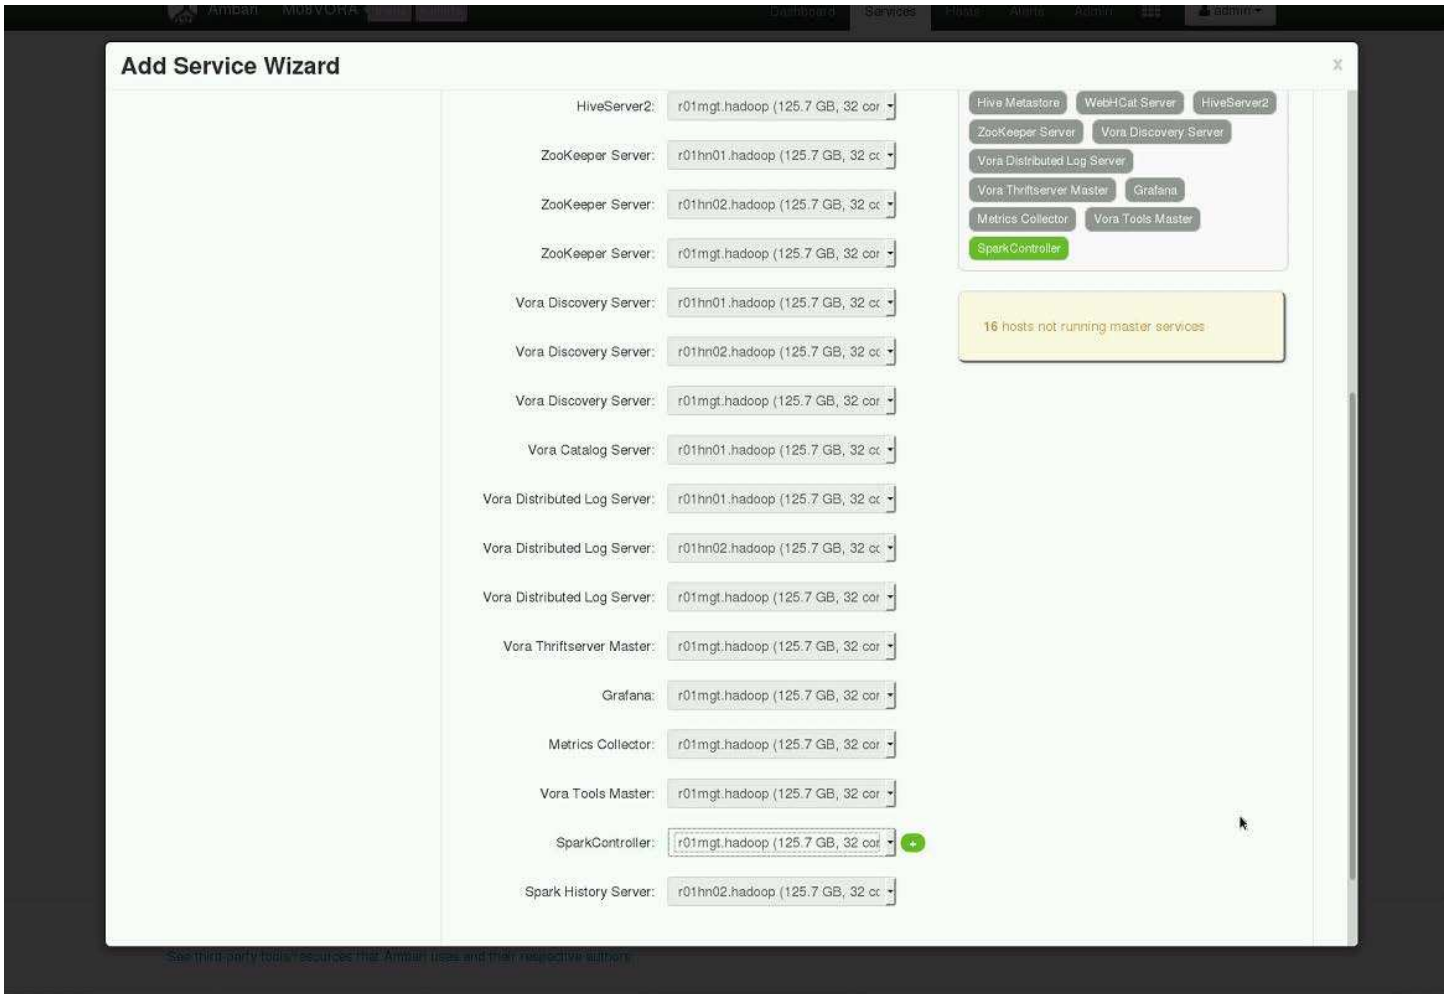
Installing SAP HANA Spark Controller

SAP HANA Spark Controller allows you to connect from SAP HANA to SAP HANA Vora and query SAP HANA Vora tables. This component is also used by tools like Data Lifecycle Manager (DLM), part of SAP HANA Data Warehousing Foundation (DWF) toolset, to enable data tiering between SAP HANA and Hadoop.

There are two ways to install the SAP HANA Spark Controller: a manual procedure and an automated way using Ambari. For the HPE EPA system, we recommend using Ambari to avoid any configuration errors. Use the following steps to install and configure the SAP HANA Spark Controller using Ambari:

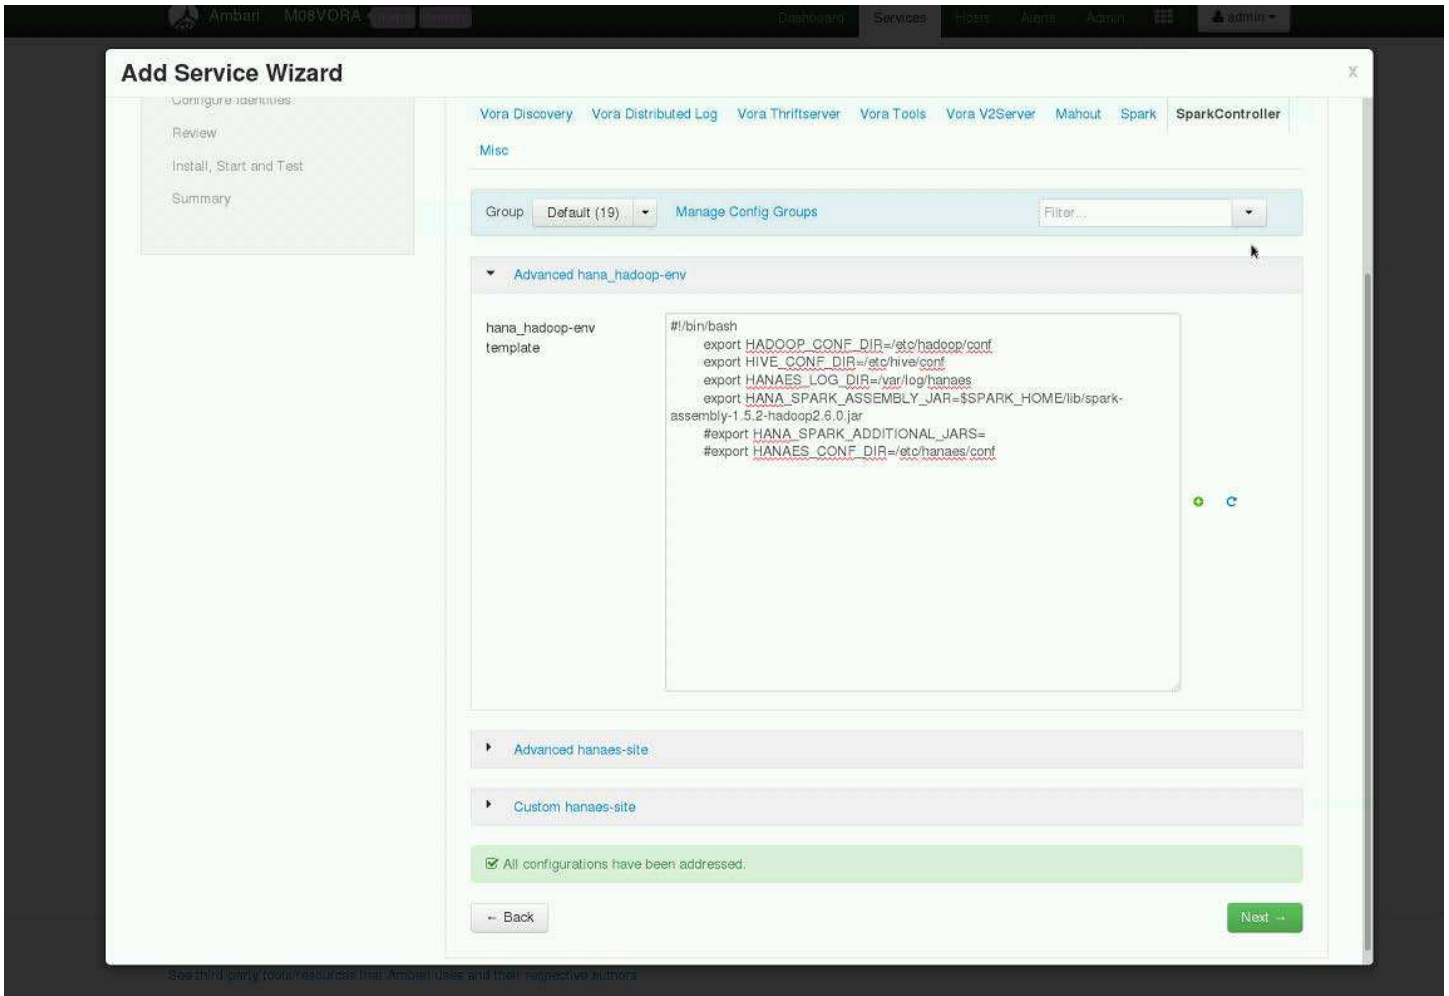
- Download the Spark Controller package from the SAP Service Market Place to the management node (**r01mgt**).
- Extract the zip file. There are two files in the zip file: a tar.gz (**controller.distribution-1.6.0-Ambari-Archivetar.gz**) file and an rpm file. Copy the tar.gz file to the Ambari Server services folder, by using the following commands:
 - cp controller.distribution-1.6.0-Ambari-Archivetar.gz /var/lib/ambari-server/resources/stack/HDP/2.3.0/services
 - tar -xvf controller.distribution-1.6.0-Ambari-Archivetar.gz

- Restart Ambari Server using the following command:
 - # ambari-server restart
- In the Ambari dashboard, select **Actions** → **Add Service** to open the **+ Add Service Wizard**.
- On the **Choose Services** screen, select the **SparkController** service.
- On the **Assign Masters** screen, assign the **SparkController** service to the management node as shown in the following figure:



See third-party tool's resources for Ambari usage and their respective authors.

- On the **Customize Services** screen, expand the **Advanced hana_hadoop-env** option, and set the properties as shown in the following figure:



See third party tools/vendors that Ambari uses and their respective authors

- On the same **Customize Services** screen, expand the **Advanced hanaes-site** option, and set the properties as shown in the following figure:

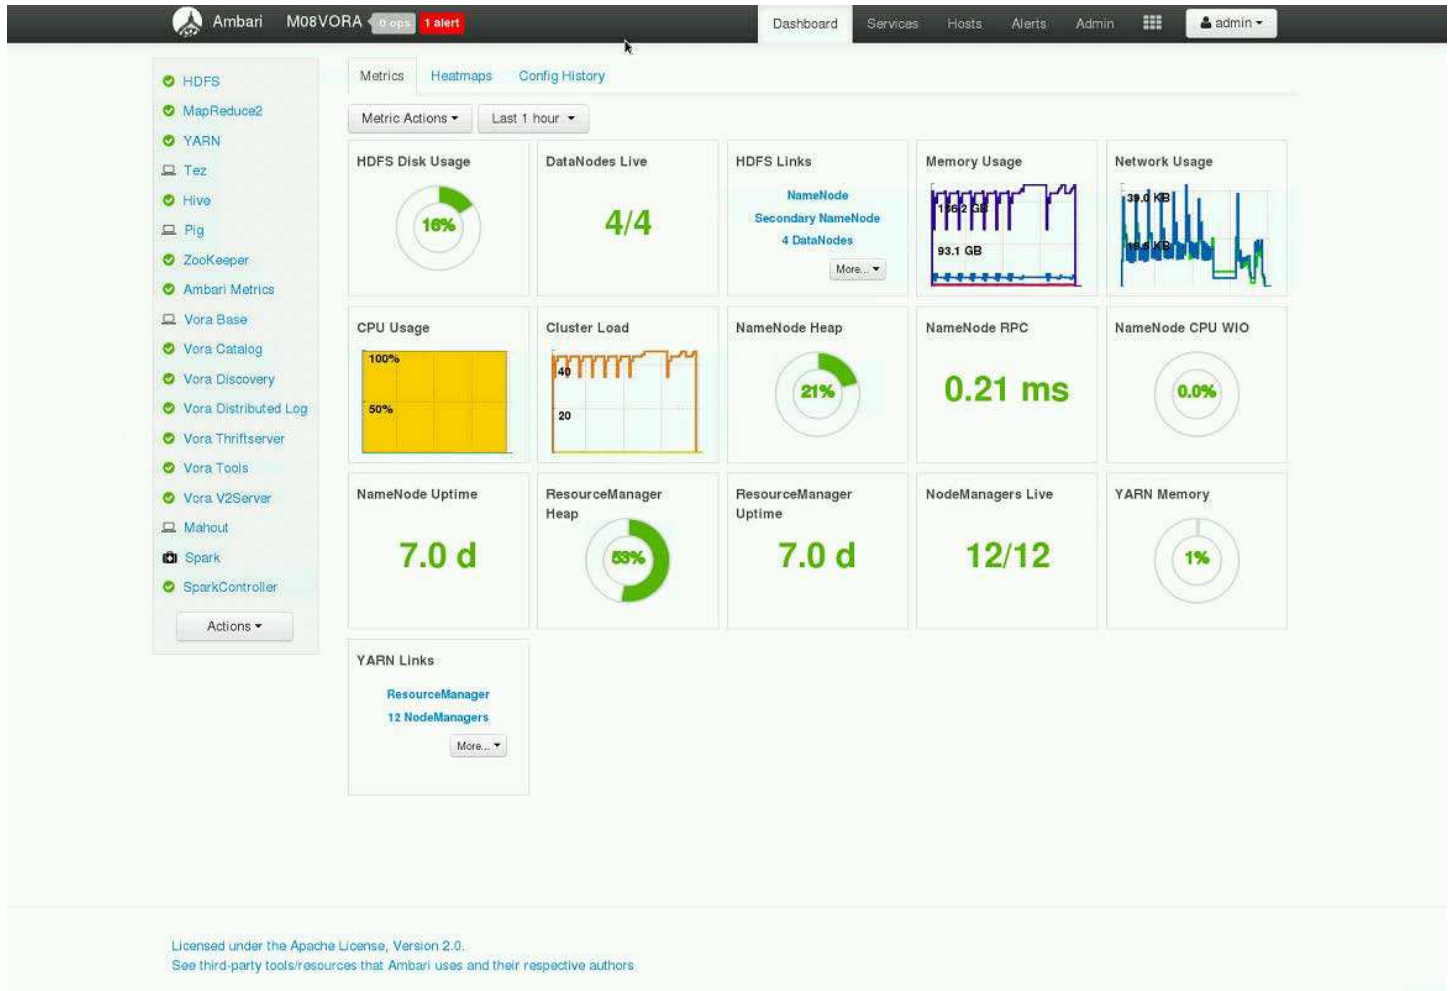
▼ **Advanced hanaes-site**

sap.hana.es.server.port	7860	○	C
sap.hana.es.spark.yarn.jar	file:///usr/sap/spark/controller/lib/external/spark-assembly-1.5.2-hadoop2.6.0.jar	○	C
sap.hana.es.ssl.clientauth.required	true	○	C
sap.hana.es.ssl.enabled	false	○	C
sap.hana.es.ssl.verify.hostname	true	○	C
sap.hana.hadoop.datastore	vora	○	C
spark.dynamicAllocation.enabled	false	○	C
spark.dynamicAllocation.maxExecutors	20	○	C
spark.dynamicAllocation.minExecutors	10	○	C
spark.executor.instances	12	○	C
spark.executor.memory	3g	○	C
spark.shuffle.service.enabled	false	○	C

▶ **Custom hanaes-site**

Verify the SAP HANA Spark Controller installation

After the SAP HANA Spark Controller service is installed, go to the Ambari Dashboard and make sure it is installed without errors and is up as shown in the following figure:



Installing SAP HANA Vora Zeppelin Interpreter

Zeppelin is a web-based graphical user interface that allows data scientists to interact easily with a Spark cluster to perform interactive data analytics. The SAP HANA Vora Spark extension provides an interpreter for the Zeppelin user interface. A modified Zeppelin interpreter is required to allow Zeppelin to run in the modified SQL context which allows SAP HANA Vora extensions. Use the following steps to install Zeppelin and modify it for SAP HANA Vora:

- Build the Zeppelin 0.5.6 package following the instructions provided on the Apache Zeppelin website (<https://zeppelin.apache.org>). Make sure that the Spark 1.5, Hadoop 2.6 and YARN profiles are selected for the build.
- Extract the built zeppelin-0.5.6-incubating.tgz archive to /home/Vora/zeppelin-0.5.6 directory as **Vora** user.
- Copy spark-sap-datasources-1.2.33-assembly.jar to /home/Vora/zeppelin-0.5.6/interpreter/spark:


```
- $ cp /var/lib/ambari-agent/cache/stacks/HDP/2.3/services/Vora-base/package/lib/Vora-spark/lib/spark-sap-datasources-1.2.33-assembly.jar /home/Vora/zeppelin-0.5.6/interpreter/spark/spark-sap-datasources-assembly.jar
```

- Combine the Zeppelin Spark interpreter JAR with the spark-sap-datasources-assembly JAR:
 - \$ cd `/home/vora/zeppelin-0.5.6/interpreter/spark`
 - \$ mkdir tmp
 - \$ [cd tmp; jar -xf ../spark-sap-datasources-1.2.33-assembly.jar]
 - \$ [cd tmp; jar -xf ../zeppelin-spark-0.5.6-incubating.jar]
 - \$ jar -cvf zeppelin-spark-sap-combined.jar -C tmp .
 - \$ rm spark-sap-datasources-1.2.33-assembly.jar
 - \$ rm zeppelin-spark-0.5.6-incubating.jar
- Add the following variables to the /home/Vora/zeppelin-0.5.6/conf/zeppelin-env.sh file:
 - export MASTER=yarn-client
 - export ZEPPELIN_PORT=9099
- Add the interpreter class `org.apache.spark.sql.SapSqlInterpreter` to the `zeppelin.interpreters` property in the `<ZEPPELIN_HOME>/conf/zeppelin-sites.xml` file.
 Make sure that the SAP interpreter class `org.apache.spark.sql.SapSqlInterpreter` occurs after the Spark interpreter class `org.apache.zeppelin.spark.SparkInterpreter` in the resulting list of interpreters.


```
...
```

```
<property>
  <name>zeppelin.interpreters</name>
  <value>INTERPRETER_1,...,INTERPRETER_N,org.apache.spark.sql.SapSqlInterpreter
</value>
  <description>Comma separated interpreter configurations.
      First interpreter becomes the default
  </description>
</property>
```
- On the Ambari Dashboard, select the YARN service and choose the **Configs → Advanced** tab. Under the **Custom yarn-site** section, click **Add Property**.
 - Add a property with the key **hdp.version** and value **2.3.0.0-2557**
- Start the Zeppelin server:
 - \$ /home/vora/zeppelin-0.5.6/bin/zeppelin-daemon.sh start
- In a web browser, navigate to <http://r01mgt:9099>. The Zeppelin Home page should load and be ready for the Vora interpreter.
- Create a new notebook by clicking the **Notebook** tab and then selecting **+ Create new note**. When you click on the **Interpreter** tab, you should see an additional interpreter prefix called **%Vora** in the Spark interpreter list.
- Enter the following script in a cell in the notebook and it should be successful, displaying an empty list of SAP HANA Vora tables:
 - %VORA SHOW TABLES

Appendix E: HPE appliances for SAP HANA details

SAP HANA, short for “High-performance Analytic Appliance,” is a data platform built for running in-memory transactional and analytical workloads. Unlike traditional databases that require these workloads to be separated, SAP HANA offers a faster and simplified architecture that consolidates both online analytic processing (OLAP) and online transaction processing (OLTP) landscapes into a single database.

HPE offers multiple solutions for SAP HANA: HPE ConvergedSystem appliance, TDI, and base configurations.

HPE ConvergedSystem 500 for SAP HANA Scale-up and Scale-out appliances

SAP HANA is SAP’s in-memory computing platform. SAP is strategically porting all their existing business applications to SAP HANA. As more and more customers are realizing the benefits of in-memory technology they are accelerating their plans to deploy/port their existing SAP applications to SAP HANA. HPE ConvergedSystem for SAP HANA configurations have been available for several years.

HPE ConvergedSystem 500 for SAP HANA Scale-up and Scale-out configurations are certified SAP HANA appliances. These are preconfigured SAP HANA appliances that act as database servers with more than a 1000 combinations of options, including memory size, number of processors, type of processors, disk configurations, operating system, and encryption option. These are Intel Xeon E7 v4 processor-based servers with SAP HANA version SPS12.

The HPE ConvergedSystem 500 for SAP HANA Scale-up configurations are available in memory sizes ranging from 128 GB to 4 TB to suit a variety of SAP HANA uses cases across analytics to SAP Business Suite on HANA. For larger scale-up configurations, the HPE ConvergedSystem 900 for SAP HANA Scale-up appliance supports up to 16 TB of memory.

External storage is available for all scale-up configurations and is based on the HPE D3700 SAS storage disk enclosure. The Small Form Factor (SFF) HPE D3700 disk enclosure with 25 drive bays offers modular solutions to simplify capacity expansion of HPE ProLiant server environments.

The front and rear of the HPE ConvergedSystem 500 for SAP HANA Scale-up appliance are shown in Figure E-1.

HPE ConvergedSystem 500 for SAP HANA Scale-up Configurations with the Intel® Xeon® E7 v4 architecture



1. HPE ConvergedSystem 500 for SAP HANA Compute Node 2. HPE ConvergedSystem 500 for SAP HANA Scale-up Storage

Figure E-1. HPE ConvergedSystem 500 for SAP HANA Scale-up appliance

Customers who are performance bound can scale-out with the modular, standards-based architecture available in configurations up to 68 TB for the HPE ConvergedSystem 500 for SAP HANA Scale-out appliance and up to 96 TB for the HPE ConvergedSystem 900 for SAP HANA Scale-out appliance.

All certified HPE appliances for SAP HANA, with Intel Xeon E7 v4 architecture, are listed at the following link:

<https://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/appliances.html#categories=Hewlett%20Packard%20Enterprise%2CIntel%20Broadwell%20EX%20E7>

The following table summarizes the HPE ConvergedSystem for SAP HANA appliances that can be connected to SAP HANA Vora.

HPE – SAP HANA based appliances	Memory for HANA
HPE ConvergedSystem 500 for SAP HANA Scale-up or Scale-out	Scale-up 128 GB to 4 TB Scale-out up to 68 TB
HPE ConvergedSystem 900 for SAP HANA Scale-up or Scale-out	Scale-up to 16 TB Scale-out to 96 TB

Appendix F: SAP HANA configuration

Setting up SAP HANA to connect to HPE EPA cluster

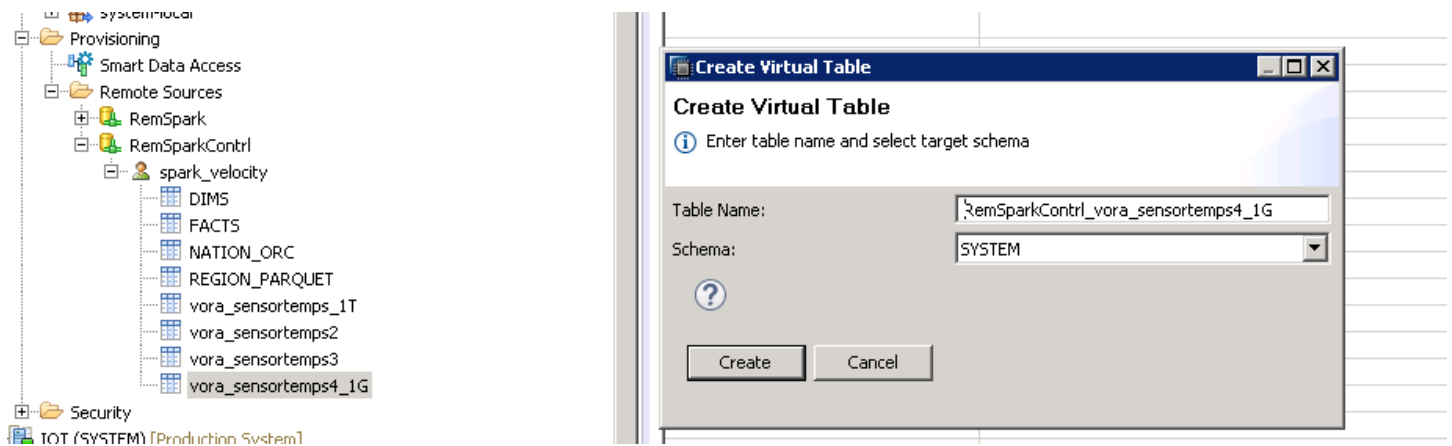
This section provides high-level instructions on configuring SAP HANA appliances to connect to the HPE EPA cluster.

Installation involves standard setup of HANA and configuration. A remote source is created using the Spark SQL destination to connect to the Big Data environment which has SAP HANA Vora running.

The detailed documentation, as listed in the “SAP HANA Spark Controller 1.6 Patch Level 0 Installation and Configuration” documentation, is to create a remote source:

```
CREATE REMOTE SOURCE "RemSparkControl" ADAPTER "sparksql" CONFIGURATION
'port=7860;ssl_mode=disabled;server=' WITH CREDENTIAL TYPE 'PASSWORD' USING
'user=hanaes;password=hanaes'
```

The diagram below shows that once the connectivity is established, tables in SAP HANA Vora can be seen in SAP HANA Studio through SAP HANA Smart Data Access. For use cases in which Big Data information needs to be accessed in SAP HANA, these tables can be exported as virtual tables and the data in SAP HANA Vora can be seen and queried in SAP HANA.



Setting up SAP HANA Data Warehousing Foundation – Data Lifecycle Manager

The Data Lifecycle Manager from SAP HANA Data Warehousing Foundation from SAP provides an easy option to relocate data to Hadoop (Spark SQL) based on aging rules. This is an additional software tool, and it has a separate license from SAP, mainly to automatically move and archive data from SAP HANA based on data aging. The details of this software configuration are available on the SAP help websites at http://help.sap.com/hana_options_dwf and http://help.sap.com/hana/SAP_HANA_DWF_Installation_Guide_en.pdf

Notes: The detailed installation is described in the SAP HANA installation guide. Some of the things to take note of are 2092669 – Release Note SAP HANA Data Warehousing Foundation.

- SAP **HANA DWF 1.0 SP03** (based on SAP HANA SP11): Install / Upgrade to SAP HANA Revision 1.00.110 or later
- SAP **HANA DWF 1.0 SP04** (based on SAP HANA SP12): Install / Upgrade to SAP HANA Revision 1.00.121, **but not beyond 1.00.121** (refer to SAP Note 2345542 for details)

- The Data Lifecycle Manager of the SAP HANA Data Warehousing Foundation 1.0 does support "Spark SQL" (Hadoop) Storage Destination starting with Support Package 03.

In addition, ensure that the Spark Controller is also at an appropriate level, as shown in the table below.

Spark Controller Version	Spark version	HANA version	DLM support	Vora connectivity
1.0 Patch Level 0	1.3.1	SP 10	No	No
1.5 Patch Level 0	1.4.1	SP 10	Limited	Yes
1.5 Patch Level 5	1.5.2	SP 11	Yes	Yes

- In the SAP HANA configuration files, data aging has to point to the remote source that was created to connect to SAP HANA Vora.

The screenshot shows a configuration file browser with a filter set to 'spark'. The table below represents the visible configuration entries:

Name	Default	System	Host - apspal46-49
global.ini			
spark_communication			
sslkeystore	sparksq_ks.pse		
ssltruststore	sparksq_ts.pse		
sslvalidatecertificate	true		
sslvalidatehostnameincertific	true		
indexserver.ini			
data_aging			
spark_remote_source		RemSparkContrl	
xsengine.ini			
data_aging			
spark_remote_source		RemSparkContrl	

- The figure below shows the storage destination configured as Spark SQL, which is needed to connect to the SAP HANA Vora system.

The screenshot shows the SAP Data Lifecycle Manager interface. The 'Storage Destinations' section is active, displaying a table with the following details for the 'RemSparkContrl' destination:

Header	
Name	RemSparkContrl
Storage Destination Type	Spark SQL (DESTINATION)
Changed By	SYSTEM
Valid From	Jul 14, 2016, 8:58:25 AM
Activation Status	Active
Version Number	1

Resources and additional links

SAP HANA Vora: <http://go.sap.com/product/data-mgmt/hana-vora-hadoop.html>

Hortonworks: hortonworks.com

Hortonworks partner site: hortonworks.com/partner/hpe

HPE Solutions for Apache Hadoop: hpe.com/us/en/solutions/big-data.html

HPE Insight Cluster Management Utility (Insight CMU): hpe.com/info/cmu

HPE FlexFabric 5930 switch series: hpe.com/networking/5930

HPE Apollo systems: hpe.com/servers/apollo

HPE ProLiant servers: hpe.com/servers/proliant

HPE Integrated Lights-Out (iLO): hpe.com/info/ilo

HPE Product Bulletin (QuickSpecs): hpe.com/info/qs

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

About Hortonworks

Hortonworks develops, distributes and supports the only 100% open source Apache Hadoop data platform. Our team comprises the largest contingent of builders and architects within the Hadoop ecosystem who represent and lead the broader enterprise requirements within these communities. The Hortonworks Data Platform provides an open platform that deeply integrates with existing IT investments and upon which enterprises can build and deploy Hadoop-based applications. Hortonworks has deep relationships with the key strategic data center partners that enable our customers to unlock the broadest opportunities from Hadoop. For more information, visit hortonworks.com



About SAP

As the market leader in enterprise application software, SAP is at the center of today's business and technology revolution. SAP helps you streamline your processes, giving you the ability to use live data to predict customer trends – live and in the moment. Across your entire business. When you run live, you run simple with SAP. SAP HANA Vora bridges the gap between corporate data and Big Data – and gives you the context you need to make better decisions. For more information, visit <http://go.sap.com/product/data-mgmt/hana-vora-hadoop.html>



Sign up for updates

© Copyright 2016 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

SAP and SAP HANA are registered trademarks of SAP AG in Germany and other countries. Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. Linux is a registered trademark of Linus Torvalds in the U.S. and other countries. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.