

Lenovo Converged Analytics Platform for SAP Vora



Easy to implement hardware, software, and services for enterprise analytics



HIGHLIGHTS

- Enable precision decisions leveraging enterprise data in SAP® HANA® and unstructured data from your Big Data environment.
- Build a system landscape that seamlessly scales and avoids platform down time.
- Have confidence your data is safe and available with MapR-FS replication – even if a node in the cluster fails.

Unlock Business Potential from your Big Data Faster and Easier

The Lenovo Converged Analytics Platform for SAP Vora is an engineered solution that extends SAP analytics across your Big Data environment. Incorporating SAP HANA, SAP Vora, and powered by the MapR Converged Data Platform, SUSE Enterprise Linux, and Lenovo systems, this solution is fast, resilient, and scalable to hundreds of nodes.

For digital businesses that want to infuse business decisions with valuable context from new data sources, SAP® Vora™ is an in-memory query engine that plugs into the Apache Spark execution framework to provide enriched interactive analytics on data stored in Hadoop. It lets you combine Big Data with corporate data in a way that is both simple and fast.

- Enable more precision decisions through greater contextual awareness in enterprise applications and analytics
- Democratize access for data scientists to facilitate new discoveries
- Simplify Big Data ownership through combined data access and processing

The MapR Platform integrates natively with SAP HANA Vora with Spark as the execution engine that enables access to data stored in the MapR Platform. It is very scalable with the ability to go to thousands of nodes and has built-in features such as High Availability and Disaster Recovery. Customers have the ability to run enhanced Spark SQL semantics seamlessly on the MapR Platform to include hierarchies that enable OLAP and drill-down analysis on large datasets.

Solution Brief

Lenovo Converged Analytics Platform for SAP Vora

The Lenovo Converged Analytics Platform for SAP Vora is an Engineered Solution designed in collaboration with SAP, MapR, SUSE, and Lenovo.



Lenovo x3650 M5 2U rack server



Lenovo x3850 X6 4U rack server



Lenovo G8272 1U 10 GbE RackSwitch

Configuration Details

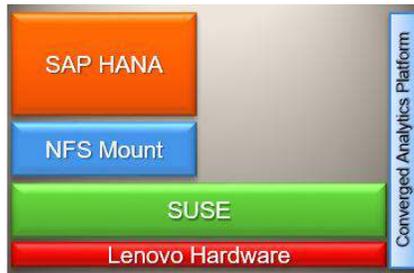
The Converged Analytic Platform for SAP Vora configuration consists of an SAP HANA system or cluster connected to a big data cluster running MapR and Hadoop. The Big Data cluster consists of Lenovo rack servers, like the powerful two-socket x3650 M5, to enhance performance and reduce power consumption of big data clusters. Purpose built for big data workloads, the 2U two-socket x3650 M5 server supports industry leading data storage capacity, the latest Intel Xeon E5-2600 v4 high performance compute processing power, flash storage options and energy efficient features. The core MapR reference architecture leverages this server as a data node for scale-out clusters.

SAP HANA runs on System x3850 X6 or x3950 X6 servers in single node or cluster configurations. These systems provide enterprise class reliability, world record SAP application performance¹ and seamless scalability for SAP HANA applications.

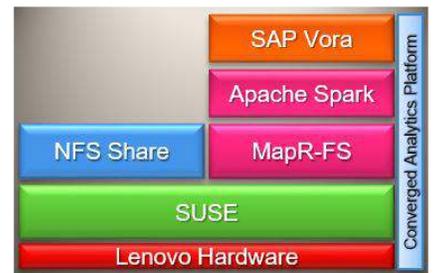
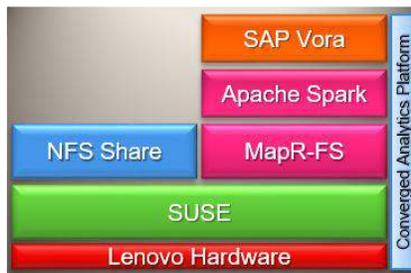
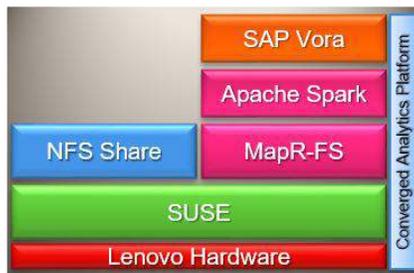
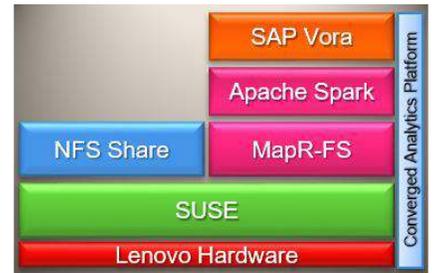
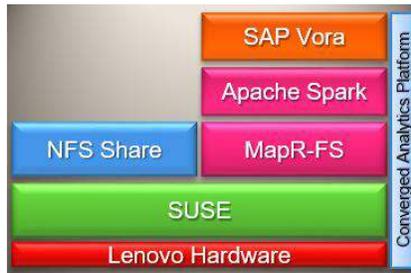
The redundant 10 Gb Ethernet data network uses the RackSwitch G8272 To connect all nodes in the cluster. Each node connects to a pair of 10 GbE switches to form a high performance and redundant active-active link

Each node in this engineered solution has internal storage that's managed by a global software defined file system that incorporates data disks from across all nodes into a single solution-wide name space - no external SAN required ! Increased capacity and performance can be achieved by seamlessly adding additional nodes and avoiding platform down time. Better still, no file system or storage changes, nor new skill sets are required of your team when scaling production environments. With MapR-FS replication, your data is safe and available, even on failure of any node in the cluster.

One or more SAP HANA nodes...



...and five or more SAP Vora nodes.



Solution Brief

Lenovo Converged Analytics Platform for SAP Vora

Technology Stack

SAP Vora is an in-memory, distributed computing solution that helps organizations uncover actionable business insights from Big Data. Use it to run enriched, interactive analytics on both enterprise and Hadoop data, quickly and easily. It runs on a Hadoop cluster and is tightly integrated with Apache Spark.

Vora enables OLAP-style capabilities with Big Data stored in Hadoop, provides deep integration with SAP HANA enabling high-performance enterprise analytics, and delivers contextual insights by combining data in SAP HANA and Hadoop.

The integration of SAP HANA with Vora on Spark makes for a database that is faster, more stable, and more scalable than any other database in the world.

SAP HANA is an in-memory, column-oriented, relational database management system with the ability to perform predictive analytics, spatial data processing, text analytics, text search, streaming analytics, and graph data processing. It also includes ETL capabilities and an application server.

The MapR Converged Data Platform enables organizations to create intelligent applications that integrate analytics with operational systems in real time.

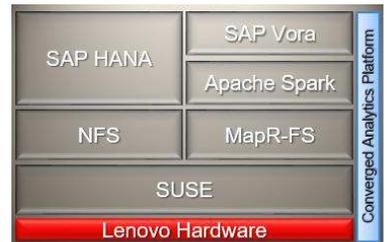
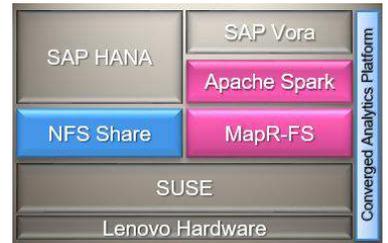
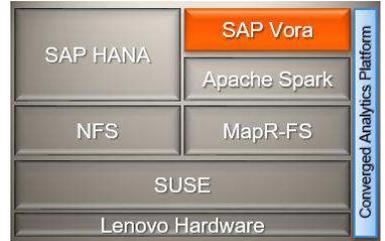
Apache Spark is a framework for blindingly fast distributed in-memory execution of code. Vora runs on top of the Spark framework.

MapR-FS is a fault tolerant, performance oriented, global filesystem that provides a single name space across all nodes in a cluster with no single point of failure. It uses a “shared nothing” model with storage distributed across all nodes in the cluster.

SUSE Linux Enterprise Server for SAP Applications provides an advanced foundation for your mission critical SAP application workloads and provides high availability with automated failover, including the SAP HANA database, operating system, and also Tailored Datacenter Integration (TDI) deployments.

For its SAP HANA Enterprise Cloud service, SAP maintains approximately 6,600 servers with over 12,000 CPUs and 16,000 virtual machines running SUSE Linux Enterprise Server and SUSE Linux Enterprise Server for SAP Applications.

Lenovo Systems and Services for SAP applications can help make your next SAP Vora deployment with MapR and SUSE a success. Lenovo supports all four leading Hadoop distributions for managing Big Data and is a leader in SAP HANA server installations, server customer satisfaction, and x86 server reliability. Lenovo has shipped over 7,200 server appliances for SAP HANA worldwide.



Solution Brief

Lenovo Converged Analytics Platform for SAP Vora



Why Lenovo

Lenovo is a leading provider of x86 servers for the data center. Featuring rack, tower, blade, dense and converged systems, and the Lenovo server portfolio provides excellent performance, reliability and security. Lenovo also offers a full range of networking, storage, software, solutions, and comprehensive services supporting business needs throughout the IT lifecycle. With options for planning, deployment, and support, Lenovo offers expertise and services needed to deliver better service-level agreements and generate greater end-user satisfaction.

For More Information

To learn more about the Lenovo Engineered Solution for the Lenovo Converged Analytics Platform, contact your Lenovo sales representative or Business Partner and visit

www.lenovo.com/sap.



© 2017 Lenovo. All rights reserved.

Availability: Offers, prices, specifications and availability may change without notice. Lenovo is not responsible for photographic or typographical errors. **Warranty:** For a copy of applicable warranties, write to: Lenovo Warranty Information, 1009 Think Place, Morrisville, NC, 27560. Lenovo makes no representation or warranty regarding third party products or services. **Trademarks:** Lenovo, the Lenovo logo, System x, ThinkServer are trademarks or registered trademarks of Lenovo. Microsoft and Windows are registered trademarks of Microsoft Corporation. Intel, the Intel logo, Xeon and Xeon Inside are registered trademarks of Intel Corporation in the U.S. and other countries. Other company, product, and service names may be trademarks or service marks of others.

CRN: BDA-LCAP-xxxx 04/2017