Solution Brief

SAP Solutions, Data & Analytics

SAP Data Hub on Lenovo and SUSE: the Servers, Containers and Storage for Better Data Orchestration

New data sources from new use cases such as IoT, video imaging, mobile devices and real-time monitoring are driving the exponential growth of data. Many enterprises are challenged to store the data as well as gain business insights and competitive advantage from it. At the same time, corporate data landscapes are becoming increasingly complex, making it difficult and costly to capture value from the available data. This solution by Lenovo and SUSE for SAP Data Hub will give your business Intelligent Insights today.

Solution at a glance:

- SUSE Containers as a Service (CaaS) offers a Kubernetes implementation to simplify orchestration and management of your solution.
- SUSE® Linux Enterprise Server for SAP Applications is the operating system that SAP itself uses for development and production.
- Lenovo ThinkSystem servers provide the high est reliability ratings among all competing x86 platforms.[1]

- **Products**

+ *SUSE Linux Enterprise Server for SAP Applications*
+ *SUSE CaaS Platform*
+ *SUSE Enterprise Storage*
+ *Lenovo ThinkSystem Server SR 530*
+ *Lenovo ThinkSystem Server SR 650*

Take Control of Your Data

Enterprises gain enormous value from data, but only if they can store, secure, manage and access it efficiently. This data orchestration is not always easy. Luckily, SAP Data Hub provides an organized and scalable approach to managing data landscapes. It can help you tie together data lakes and other vital business sources such as SAP S/4HANA through a single management console. With the Lenovo and SUSE solution for SAP Data Hub, you get an integrated, end-to-end solution built on proven components, with a flexible building-block approach that guarantees high scalability.

With the Lenovo and SUSE solution for SAP Data Hub, you can:

- Easily scale from proof of concept to production.
- Access better availability with built-in best practices from SAP, SUSE and Lenovo.
- Achieve a faster time to value with easy implementation and a flexible deployment approach.

The Challenges of Data

The modern enterprise faces a jungle of data sources. You might have critical data residing in cloud storage, SAP systems, Hadoop and more. Accessing all of this data and maintaining control of it with governance and security is a complex challenge.

Managing and processing data should be an opportunity to unlock its value and deliver new insights;—enterprises just need the right tools. SAP Data Hub delivers a

1 *"ITIC 2017 Global Reliability Survey Mid-Year Update," ITIC, June 14, 2017.*

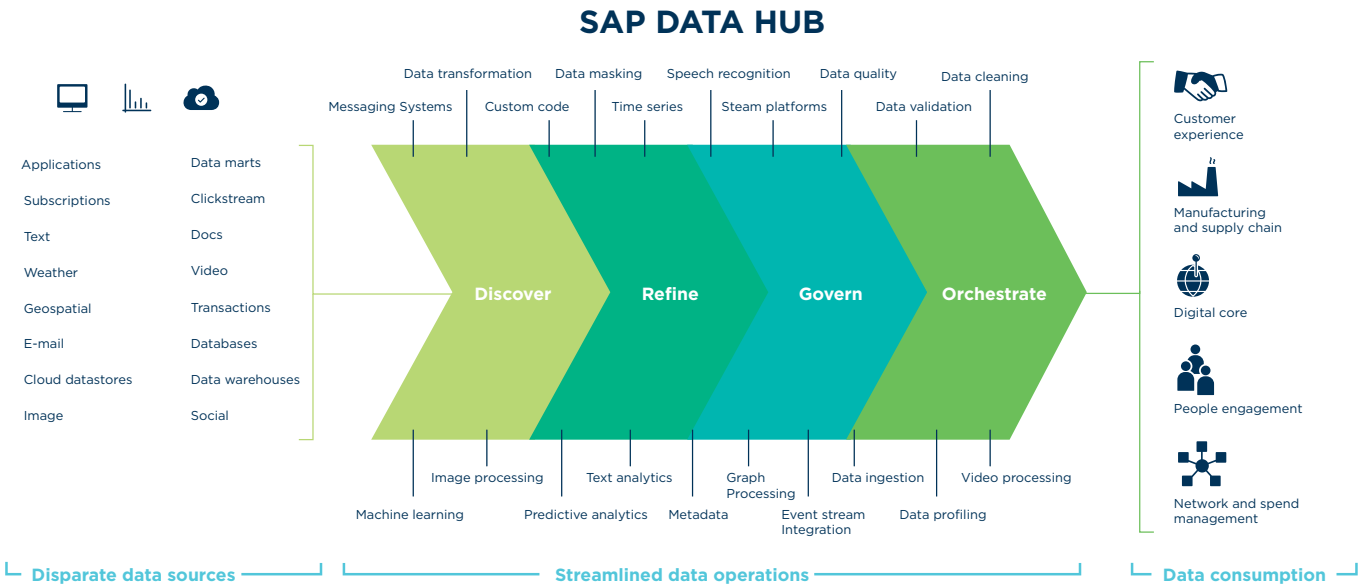

Figure 1. *End-to-End Data Orchestration*

better way to achieve that end. It is the umbrella that organizes and integrates all of the solutions, data and processing involved in SAP's overall approach to the intelligent enterprise.

SAP Data Hub provides a single data management pane for data from various sources, such as Hadoop, cloud storage, SAP HANA, business applications and more. It can help you resolve data management issues and streamline access to data.

SAP Data Hub has a number of prerequisites. You must set up SAP HANA, select a Kubernetes distribution, build a storage and networking infrastructure and select operating systems and hardware for it all. Lenovo has worked with SUSE and SAP to simplify all of this by building a robust and easy-to-deploy solution for SAP Data Hub, complete with a reference architecture.

The Solution

The Lenovo and SUSE architecture offers a fast, flexible way to support SAP Data Hub. It has been validated by Lenovo, SUSE and SAP and implements the best practices brought by each of these leading organizations.

The solution involves running SAP Data Hub as a containerized workload on SUSE CaaS:

- A secure Docker registry provides the container images.
- An optional Hadoop cluster can be connected to the system.
- A software-defined storage solution based on Lenovo and SUSE Enterprise Storage provides a reliable and scalable storage layer for the entire
- Lenovo ThinkSystem servers provide compute and storage for the solution, with networking provided by Lenovo RackSwitch.

The Hardware Foundationesti

The Lenovo ThinkSystem SR530 server is an ideal two-socket 1U rack server for businesses that need versatility; industry-leading reliability, management and security; and cost-optimized performance and flexibility for future growth. Featuring the Intel Xeon Scalable processor family, the SR530 server offers a balance of performance, capacity and value, making it an ideal foundation for the SAP Data Hub solution.

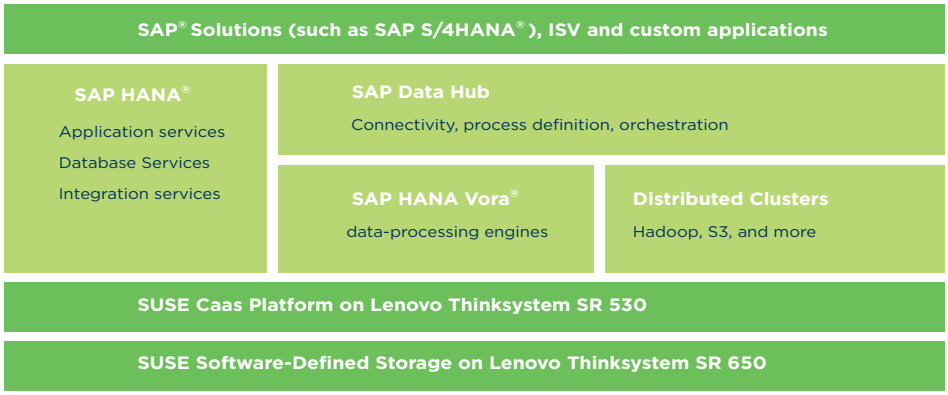

Figure 2. *Lenovo Solution for SAP Data Hub Architecture*

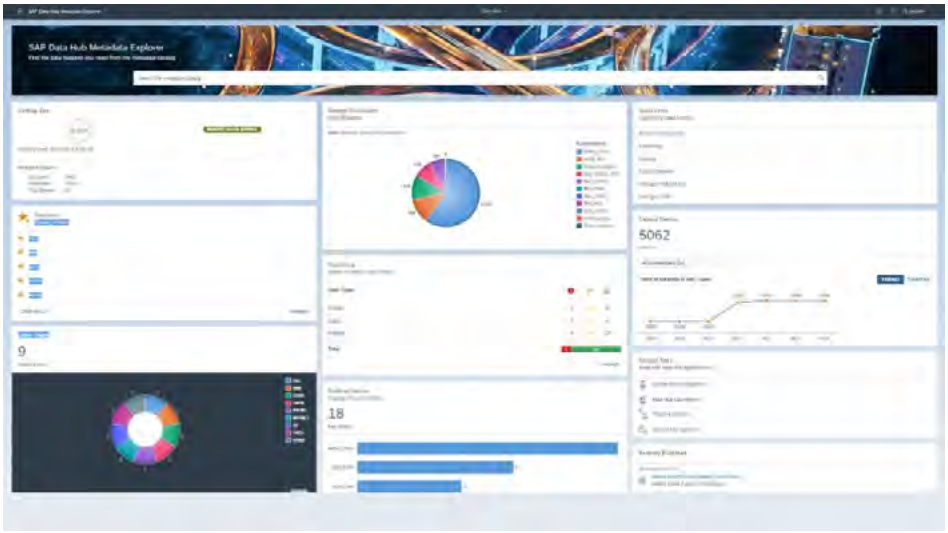

Figure 3. *SAP Data Hub Discovery*

Combining performance and flexibility, the Lenovo ThinkSystem SR650 server is the foundation for the Lenovo DSS-C software-defined storage solution. The SR650 server provides extensive storage capabilities with up to 3 TB of 2666 MHz TruDDR4 memory and up to 24x 2.5-inch or 14x 3.5-inch drive bays, with a choice of NVMe PCIe SSDs, SAS/SATA SSDs and SAS/SATA HDDs.

It can provide outstanding uptime to keep business-critical applications and cloud deployments running safely. Ease of use and comprehensive systems management tools help to make deployment easier. Outstanding reliability, availability and serviceability (RAS), as well as high-efficiency design, improve your business environment and help to save operational costs.

The Software-Defined Infrastructure

Three SUSE components provide the software-defined infrastructure for SAP Data Hub. SUSE Linux Enterprise for SAP Applications is the foundation for SAP HANA. It is the development platform for SAP and the most-used operating system for SAP HANA. The SAP-specific operating system from SUSE offers features not found anywhere else, including built-in high availability, page cache management to maintain system performance and resource agents to help automate failover in high availability scenarios. SUSE offers priority support via a single point of contact through SAP.

The second component is the SUSE CaaS platform. SUSE CaaS is an application development and hosting platform for container-based applications and services. It enables you to provision, manage and scale container-based applications and services. The easy-to-use container platform includes everything you need in a single solution: Kubernetes for orchestration, microservices and container OS, and configuration.

SUSE Enterprise Storage, powered by Ceph—a highly scalable and resilient software-based storage solution—enables organizations to build cost-efficient and highly scalable storage using off-the-shelf

Many organizations claim strong partnerships, but how many rely on one another's solutions? SAP itself relies on Lenovo and SUSE in its production and development environments, and both Lenovo and SAP are SUSE customers. The three organizations have worked closely on joint solutions and mutually supported offerings for many years.

Contact us at:
www.suse.com/lenovo
lenovo@suse.com

servers and disk drives. It is self-managing and delivers storage functionality comparable to mid- and high-end storage products at a fraction of the cost—including file, block and object storage in a single solution. In the SAP Data Hub solution, SUSE Enterprise Storage provides storage for SAP HANA and persistent volumes for containers running on SUSE CaaS.

Reference Architecture to deploy SAP Data Hub on Lenovo ThinkSystem Servers and SUSE software available.[2]

You can also add a Hadoop data lake collocated on the Kubernetes cluster, with storage on the Lenovo DSS-C Software Defined Storage solution. This enables SAP Data Hub to use the Hadoop cluster as a computational framework for SAP Data Hub jobs.

Solution Components

Five nodes (SUSE CaaS), one node (Docker registry), three optional nodes (Hadoop)

- Lenovo ThinkSystem SR530 server in the following configuration:
  - 1x Intel Xeon Silver 4114 CPU @ 2.20 GHz, 10-core
  - 192 GB (6x 32 GB DIMMs) DDR4 RAM
  - 2x 240 GB 6 Gbps SATA 2.5" SSDs in a RAID1 configuration for the OS

Four nodes (SUSE Enterprise Storage)

- Lenovo ThinkSystem SR650 server, in the following configuration:
  - 2x Intel Xeon Silver 4110 CPU @ 2.10 GHz, 8-core
  - 92 GB (6x 32 GB DIMMs) DDR4 RAM
  - 2x 240 GB 6 Gbps SATA 2.5" SSDs in a RAID1 configuration for the OS
  - Four or more 960 GB 6 Gbps SATA 2.5" SSDs (no RAID) for data

Lenovo ThinkSystem NE2572 RackSwitch
Lenovo RackSwitch G7052
Lenovo Professional Services, for installation and configuration of the solution

SUSE Linux Enterprise Server for SAP Applications
SUSE CaaS Platform
SUSE Enterprise Storage

Why Lenovo and SUSE

Lenovo is the leading provider of x86 systems for the data center. The portfolio includes rack, tower, blade, dense and converged systems. It supports enterprise-class performance, reliability and security. Lenovo also offers a full range of networking, storage, software and solutions, as well as comprehensive services that support business needs throughout the IT lifecycle.

SUSE, the world's largest independent open source provider, has a 20 year partnership with SAP to drive co-innovation so customers gain and sustain competitive advantages in the intelligent era. SUSE Linux Enterprise for SAP applications is the leading platform for SAP applications on Linux and SUSE CaaS Platform 3 on premise is validated for SAP Data Hub 2.5 with SUSE Enterprise Storage as storage backend.

2 *https://www.suse.com/media/white-paper/Reference_Architecture_for_SAP_Data_Hub_on_Lenovo_with_SUSE_CaaS_SES.pdf*