



SUSE Harvester and Rancher Solutions on ThinkSystem V3 Servers

Solution Brief

Business Trends

Cloud-Native Hyperconverged Architecture: The growth of technology, virtualization, containerization, cloud computing, artificial intelligence, and data-driven ecosystems creates a need for on-premises infrastructure to meet high-performance and high-density workload requirements. Servers need software-defined solutions to build robust local storage and integrated software stacks for cloud and cloud-native technologies, as well as third-party software, to create flexible and scalable architectures capable of meeting any workload. This trend means businesses need a common hyperconverged solution with faster servers and storage to reduce datacenter footprint and achieve the performance and low-latency objectives for both virtualization and container workloads.

Managing Virtualization and Cloud-Native Workloads : The adoption of containers, hybrid cloud, artificial intelligence and machine learning, and open-source technologies drives the need for virtual machines and container applications to coexist on shared infrastructure. Ever-growing data and AI/ML workloads require storage systems with low-latency, high-capacity drives and high-speed network adapters and interconnects to meet performance requirements. IT teams require an enterprise-grade, open-source hyperconverged solution to reduce costs and simplify deployment and management without compromising operational efficiency. Businesses need to consolidate enterprise workloads and rapidly provision end-to-end infrastructure and software stacks without compromising performance and integration capabilities.

SUSE Harvester and Rancher Solutions on ThinkSystem V3 Servers

Lenovo ThinkSystem V3 systems are the perfect choice for hyperconverged infrastructure and provide an outstanding platform to support the different workloads. Lenovo ThinkSystem V3 servers are equipped with 5th Gen Intel® Xeon® Scalable processors and 5th Gen AMD EPYC processors.

SUSE Rancher is a Kubernetes management platform which provides a single-pane-of-glass approach for provisioning, upgrading, securing, and monitoring Kubernetes clusters across on-premises and in the cloud infrastructure. Rancher is integrated with Rancher Kubernetes Engine Government (RKE2) distribution which uses containerd as an embedded container runtime.

SUSE Harvester is an open source hyperconverged infrastructure (HCI) solution designed for modern cloud native environments and it uses RKE2 as a container orchestrator on top of SUSE Linux Enterprise Server and KubeVirt technology to provide virtualization and Longhorn for distributed block storage. This architecture helps operators consolidate and simplify to run their virtual machine workloads alongside containerized workloads using Kubernetes control plane. Storage is distributed along the Harvester nodes and Longhorn scales easily with the rest of the cluster as new nodes are added and supports backups to NFS or S3 buckets.

Harvester is fully integrated with Rancher and leverages it for authentication, authorization, and cluster management and can be managed and monitored from Rancher console. This integration unifies management across multiple clusters and environments, simplifying operations, improving governance, and enhancing overall visibility and control. The hyperconverged infrastructure designed with Lenovo ThinkSystem Servers and SUSE Harvester eases and scales deployment and operations. With Rancher Prime and Harvester, IT operators now have access to an enterprise-ready, simple-to-use infrastructure platform that cohesively manages their virtual machines and Kubernetes clusters alongside one another.

Lenovo ThinkSystem V3 Servers

Lenovo ThinkSystem SR650 V3 (2U) and SR630 V3(1U) systems are equipped with 5th Gen Intel® Xeon® Scalable processors. The 5th Gen Intel Xeon Scalable processors support up to 64 cores, up to 3.9 GHz clock speed, up to 8 TB memory and support PCIe 5.0 and NVMe drives.

Lenovo ThinkSystem SR635 V3 (1U 1P), SR645 V3 (1U, 2P), SR655 V3 (2U, 1P) and SR665 V3 (2U, 2P) systems are featured with 5th generation AMD EPYC™ processors. The AMD EPYC 9005 family of processors offer up to 160 cores, up to 4.0 GHz clock speed, up to 3TB memory, and support PCIe 5.0 and NVMe drives.

SUSE Linux Enterprise Server is certified and supported on ThinkSystem V3 servers.

Applications where the servers would excel include:

- Mission Critical applications
- ERP
- CRM
- Business Intelligence (BI)
- Databases and Data Warehouse
- Analytics
- Virtual Desktops
- High Performance Computing (HPC)
- Artificial Intelligence
- Server Consolidation
- Virtualization

Longhorn Software Defined Storage Architecture

Longhorn is a thin-provisioned crash-consistent block storage system, and it has two components.

The **Longhorn Manager** Pod runs on each node in the Longhorn cluster as a Kubernetes DaemonSet. It is responsible for creating and managing volumes in the Kubernetes cluster, and handles the API calls from the UI or the volume plugins for Kubernetes.

The Longhorn Engine is a dedicated storage controller for volume created through Longhorn Manager. It creates replicas distributed across hosts in the cluster to ensure maximum availability. The multiple data paths of the replicas ensure high availability of the Longhorn volume. The Longhorn Engine always runs in the same node as the Pod that uses the Longhorn volume.

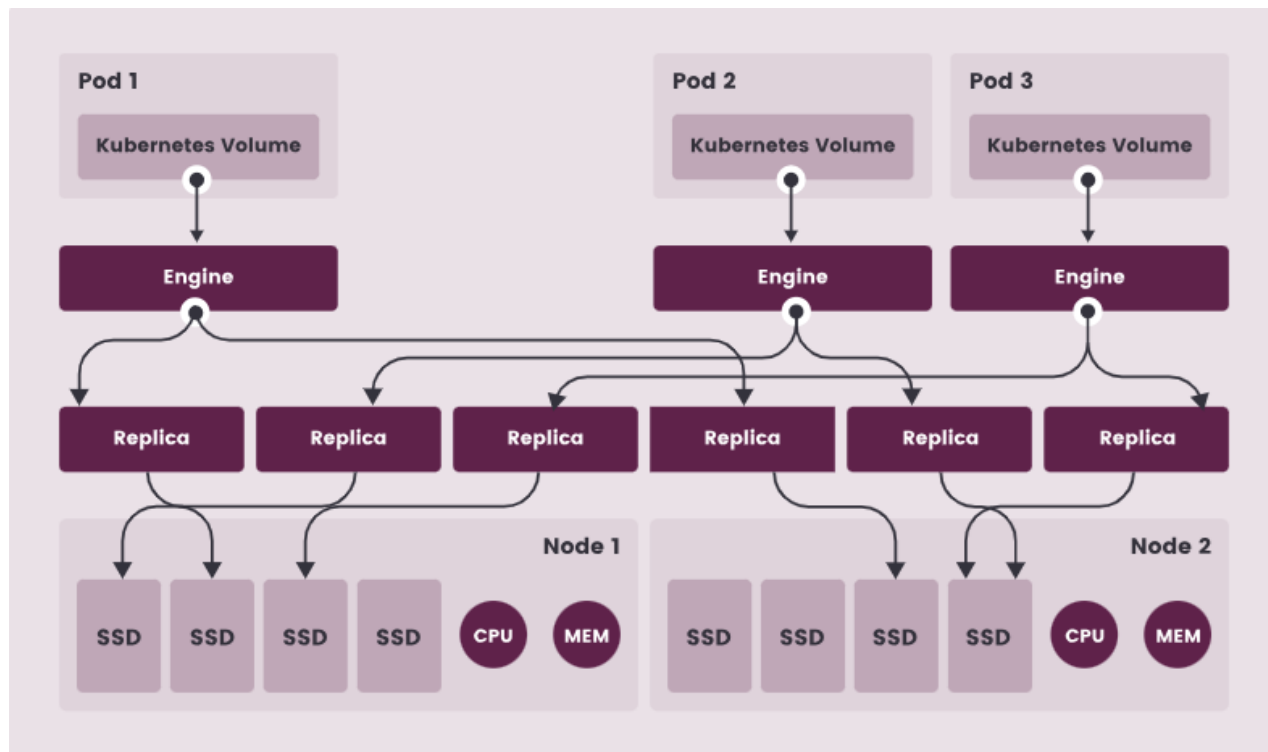


Figure 1. The Longhorn Engine

Harvester Architecture and Deployment

Harvester can start in a single node (also known as all-in-one), where the node serves as a compute and a single node control plane. Longhorn is deployed as part of Harvester, provides the storage layer. When the cluster reaches three nodes, Harvester will reconfigure itself to provide High Availability features without disruption. Harvester only supports VLAN through bridges and NIC bonding.

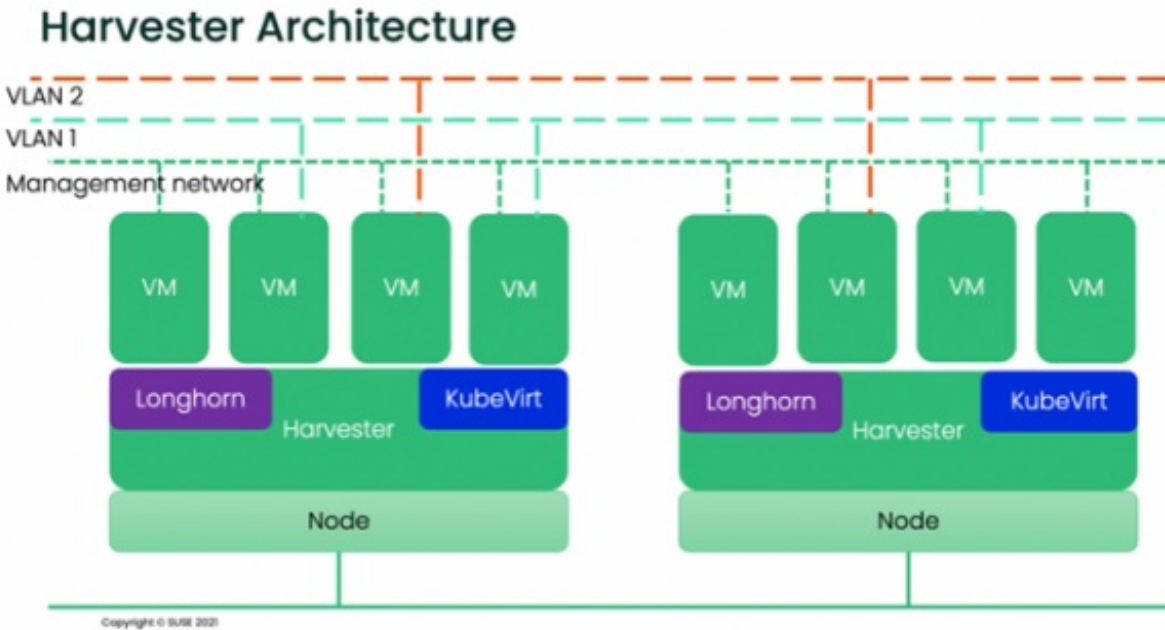


Figure 2. Harvester Architecture

Unified Management for Kubernetes Clusters and Hyperconverged Storage

Figure 3 below shows multi-cluster registration option in Harvester management console.

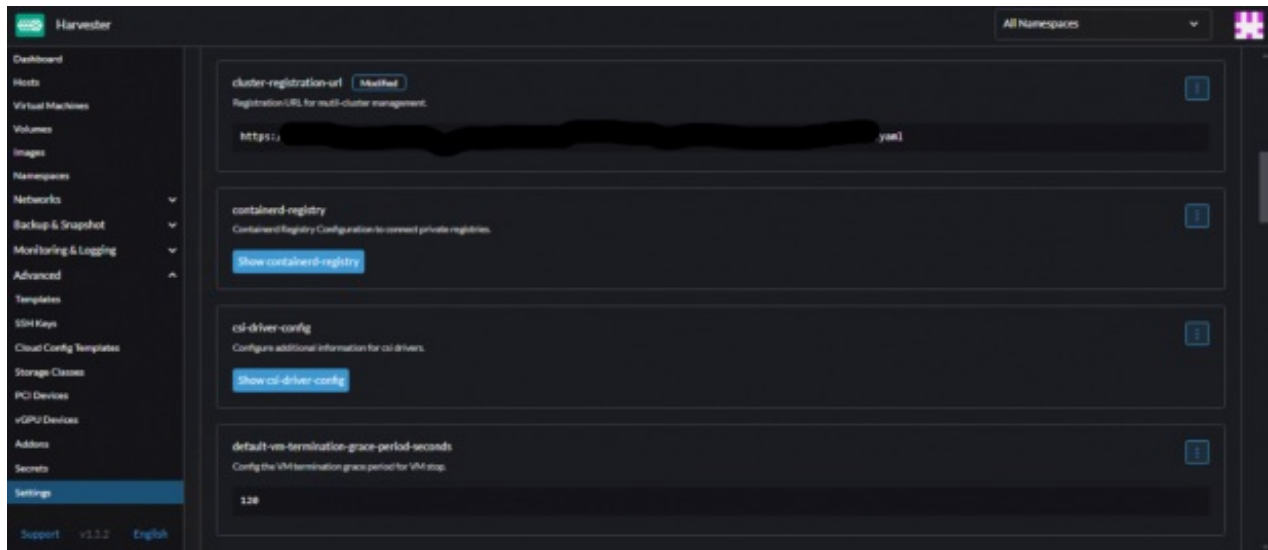


Figure 3. Harvester Management Console

Figure 4 below shows Rancher dashboard with Harvester cluster management.

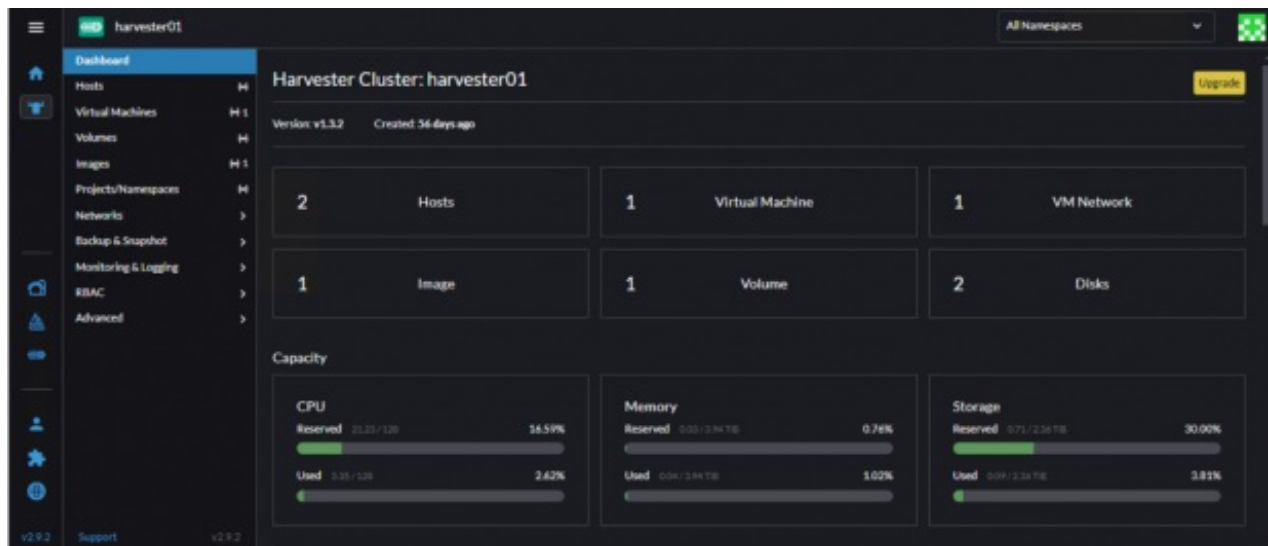


Figure 4. Rancher Dashboard

Rancher Server Deployment Model

Rancher Server includes all the software components used to manage the entire Rancher deployment. For the best performance and security, we recommend a dedicated Kubernetes cluster for the Rancher management server. Rancher can run on a single node, or on a high-availability Kubernetes cluster and A high-availability Kubernetes installation is recommended for production.

The Rancher server, regardless of the installation method, should always run on nodes that are separate from the downstream user clusters that it manages. If Rancher is installed on a high-availability Kubernetes cluster, it should run on a separate cluster from the cluster(s) it manages.

A user cluster is a downstream Kubernetes cluster that runs your apps and services. Figure 5 below depicts two downstream user clusters are managed by Rancher management server.

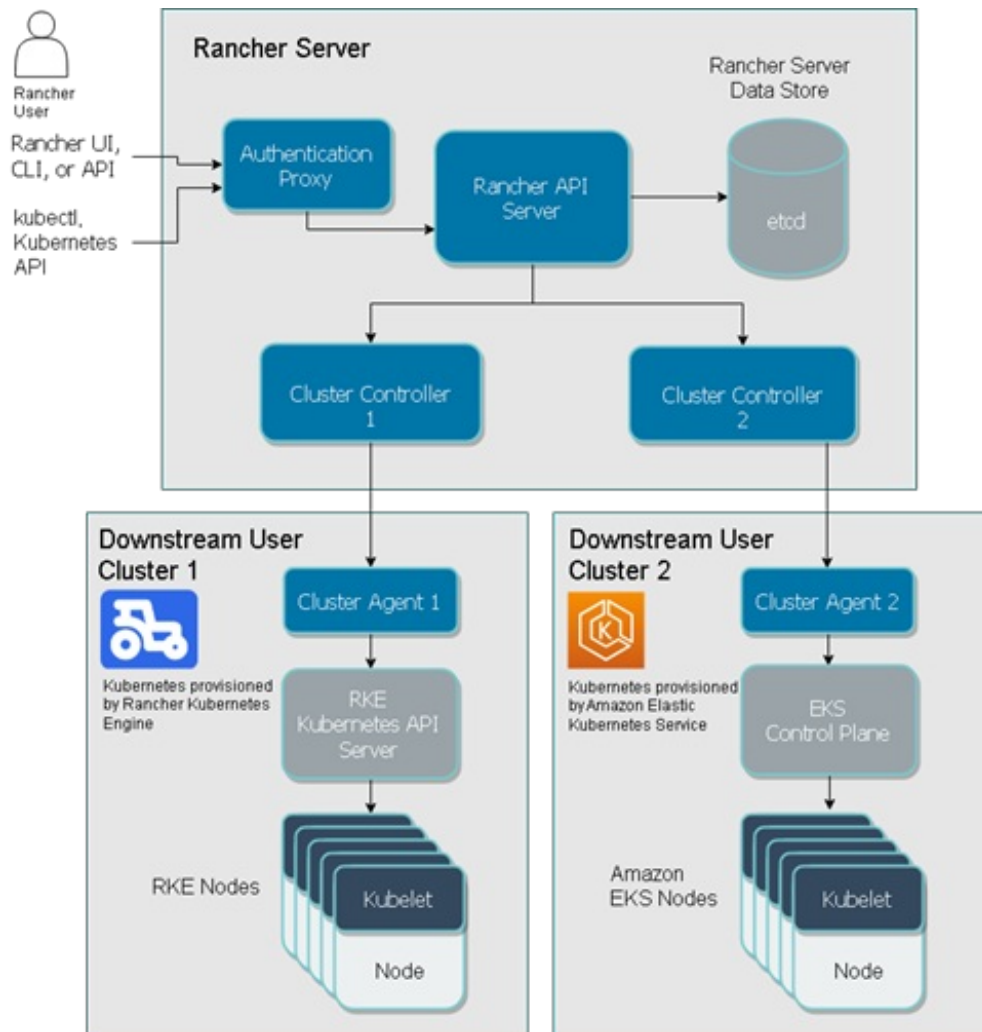


Figure 5. Rancher Management Server

Related product families

Product families related to this document are the following:

- [ThinkSystem SR630 V3 Server](#)
- [ThinkSystem SR635 V3 Server](#)
- [ThinkSystem SR645 V3 Server](#)
- [ThinkSystem SR650 V3 Server](#)
- [ThinkSystem SR655 V3 Server](#)
- [ThinkSystem SR665 V3 Server](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2025. All rights reserved.

This document, LP2110, was created or updated on January 6, 2025.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP2110>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP2110>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ThinkSystem®

The following terms are trademarks of other companies:

AMD and AMD EPYC™ are trademarks of Advanced Micro Devices, Inc.

Intel® and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Other company, product, or service names may be trademarks or service marks of others.