Reference Architecture for SUSE Rancher Prime on Lenovo ThinkSystem V4 Servers

Last update: **13 January 2026**
Version: **1.0**

Enterprise-Grade Kubernetes on Lenovo Bare-Metal

Unified Multi-Cluster and Hybrid Cloud Management

Edge Computing Enablement with Lightweight K3s

Integrated Security, Storage, and Networking for Modern Workloads

Sorin Nicolae Renghea
Maria Daniela Albu

LENOVO PRESS

Table of Contents

1 Introduction

As organizations accelerate their digital transformation, the shift toward containerized workloads and cloud-native applications has become a strategic imperative. These technologies offer agility, scalability, and portability enabling faster development cycles, simplified operations, and consistent deployment across diverse environments.

At the heart of this transformation is Kubernetes, the industry-standard platform for container orchestration. It provides robust capabilities for automating deployment, scaling, and management of containerized applications. However, managing Kubernetes at scale, especially across hybrid cloud and edge environments can be complex and resource-intensive.

To address these challenges, SUSE Rancher Prime, offers a comprehensive Kubernetes management solution. SUSE Rancher Prime simplifies the lifecycle of Kubernetes clusters by providing centralized control, policy enforcement, identity management, and seamless integration with cloud and edge platforms.

It supports both upstream Kubernetes and lightweight distributions like K3s, making it ideal for a wide range of use cases, from enterprise data centers to remote edge locations.

This reference architecture outlines a blueprint for deploying Kubernetes clusters with SUSE Rancher Prime on **Lenovo bare-metal servers**. Lenovo's **ThinkSystem** and **ThinkEdge** platforms deliver high-performance, scalable infrastructure purpose-built for container workloads. By leveraging bare-metal deployments, organizations can eliminate hypervisor overhead, gain direct access to hardware resources, and achieve superior performance and control.

By combining Rancher's advanced Kubernetes management capabilities with Lenovo's robust infrastructure, this architecture provides a solid foundation for building cloud-native environments that meet stringent requirements for **performance, scalability, security, and operational simplicity**. Whether deployed in centralized data centers, at distributed edge sites, or in hybrid scenarios, this solution empowers organizations to maximize the value of their container and cloud-native strategies.

Ultimately, this reference architecture is designed to serve as both a practical guide and a strategic framework for enterprises seeking to modernize their IT operations. By following the blueprint outlined in this document, organizations can accelerate the adoption of Kubernetes with confidence, simplify the management of complex environments, and unlock new opportunities for innovation at scale. Together, SUSE Rancher and Lenovo enable IT teams to move beyond experimentation toward building resilient, production-ready cloud-native platforms that are future-proof, efficient, and adaptable to the evolving demands of the digital landscape.

1.1 Purpose of This Document

This reference architecture provides a validated, modular, and adaptable framework for deploying Kubernetes clusters using SUSE Rancher Prime on Lenovo bare-metal infrastructure. It is designed to guide IT architects, DevOps teams, and infrastructure engineers in building scalable, secure, and high-performance container platforms that meet enterprise-grade requirements.

The document outlines best practices, supported configurations, and deployment models for both centralized data center environments and distributed edge scenarios. It aims to simplify decision making around hardware selection, software stack integration, and operational management while ensuring compliance, resilience, and future scalability.

In addition to providing technical guidance, this document serves as a reference point for aligning infrastructure design with broader organizational goals such as operational efficiency, cost optimization, and business agility. By combining architectural principles with practical deployment considerations, it helps organizations reduce complexity, accelerate time-to-value, and establish a consistent foundation for innovation across hybrid and edge environments.

1.2 Overview of SUSE Rancher + Kubernetes on Lenovo Infrastructure

SUSE Rancher Prime is a comprehensive Kubernetes management platform that simplifies multi-cluster operations across hybrid cloud and edge environments. It provides centralized authentication, policy enforcement, and lifecycle management for Kubernetes clusters, whether they are running in the cloud, on-premises, or at the edge.

Lenovo's bare-metal servers, such as the **ThinkSystem SR630/SR650 V4** for data centers and **ThinkEdge SE100/SE455 V3** for edge deployments, offer high-performance, scalable infrastructure optimized for container workloads. By eliminating hypervisor overhead and leveraging direct access to hardware, these platforms deliver superior performance, security, and control.

Together, SUSE Rancher and Lenovo provide an open, enterprise-ready solution stack that supports:

- **Kubernetes orchestration** with full CNCF conformance
- **Multi-cluster and hybrid cloud management**
- **Edge computing with lightweight K3s clusters**
- **Enterprise-grade security** including RBAC, CIS compliance, and runtime protection

This reference architecture ensures organizations can modernize their infrastructure while maintaining flexibility, performance, and operational efficiency.

2 Architecture Goals

The reference architecture is designed to provide organizations with a clear and validated pathway for deploying and managing Kubernetes clusters at enterprise scale. It emphasizes not only the technical requirements of a modern container platform but also the operational and business outcomes that enterprises expect from digital transformation initiatives. The following key goals form the foundation of this architecture:

2.1 Production-Grade Kubernetes Deployments

At its core, the architecture aims to deliver Kubernetes environments that are robust, resilient, and ready for mission-critical workloads. This includes:

- **High availability and fault tolerance**, ensuring minimal downtime and business continuity.
- **Scalability to support diverse workloads**, ranging from stateful applications to microservices-driven architectures.
- **Performance optimization**, leveraging Lenovo bare-metal servers to maximize resource efficiency and application throughput.
- **Operational reliability**, achieved through standardized deployment patterns, best practices, and automation.
  By meeting these requirements, organizations can confidently adopt Kubernetes as a production-ready platform that supports both current and future business needs.

2.2 Multi-Cluster and Hybrid Cloud Management

Modern enterprises rarely operate within a single environment. This architecture is designed to simplify management across multiple Kubernetes clusters and heterogeneous infrastructures. SUSE Rancher Prime plays a central role by providing:

- **Centralized control and visibility** across clusters, regardless of whether they run in on-premises data centers, public clouds, or edge locations.
- **Consistent policy enforcement and governance**, reducing operational risk and ensuring compliance across diverse environments.
- **Seamless workload portability**, enabling applications to move between environments without refactoring.
- **Support for hybrid and multi-cloud strategies**, giving organizations the flexibility to optimize costs, performance, and resilience by leveraging the right platform for each workload.
  This unified approach ensures that IT teams can manage complexity at scale without sacrificing agility or control.

2.3 Edge Computing Enablement

As digital services extend closer to end-users, devices, and data sources, edge computing has become a critical component of enterprise IT strategies. This architecture supports edge enablement through:

- **Lightweight Kubernetes distributions (K3s)**, optimized for resource-constrained environments.
- **Lenovo ThinkEdge platforms**, providing ruggedized, high-performance hardware purpose-built for edge deployments.
- **Centralized management of distributed edge clusters**, ensuring consistent operations across thousands of sites.
- **Support for real-time processing and low-latency workloads**, allowing organizations to deliver responsive services where data is generated.
  With these capabilities, enterprises can expand their digital footprint beyond traditional data centers and bring intelligence to the edge while maintaining consistent security and governance.

2.4 Enterprise-Grade Security and Compliance

Security is integral to every aspect of Kubernetes adoption, from infrastructure to application deployment. This reference architecture embeds enterprise-grade security and compliance by design:

- **Identity and access management (IAM)** integration, ensuring secure and role-based cluster operations.
- **Policy-driven security controls**, including admission control, network segmentation, and workload isolation.
- **Continuous compliance monitoring**, delivered through SUSE Rancher Prime solutions, enabling organizations to adhere to regulatory requirements such as GDPR, HIPAA, or PCI DSS.
- **End-to-end encryption and data protection**, leveraging Kubernetes-native features (such as TLS for API server communication and secrets encryption at rest) combined with SUSE Rancher Prime and SUSE Security capabilities for securing data in transit and at rest across data center, cloud, and edge environments.

By addressing security holistically, the architecture ensures that Kubernetes clusters remain compliant, resilient against threats, and aligned with enterprise risk management strategies.

This framework of goals establishes the guiding principles for designing and deploying Kubernetes environments with SUSE Rancher Prime on Lenovo bare-metal infrastructure. Together, these objectives ensure that organizations can build platforms that are **resilient, flexible, secure, and future-ready**.

3 Architectural Overview

SUSE Rancher Prime provides a comprehensive platform for deploying and managing Kubernetes clusters across on-premises, cloud, and edge environments. It includes integrated tools for application deployment, monitoring, and security, while streamlining multi-cluster operations through a unified interface. SUSE Rancher Prime supports leading Linux distributions including SUSE Linux Enterprise Server (SLES), Ubuntu, RHEL, and is compatible with modern container runtimes such as containerd and CRI-O. Built-in access controls and identity integrations ensure secure, policy-driven cluster management.

Our solution uses the following:

- Bare-metal environment: Lenovo ThinkSystem SR630 V4 or Lenovo ThinkSystem SR650 V4
- Operating system: SUSE Linux Micro
- Container application platform: SUSE Rancher Prime Kubernetes Engine 2 (RKE2)
- Persistent storage: SUSE Longhorn

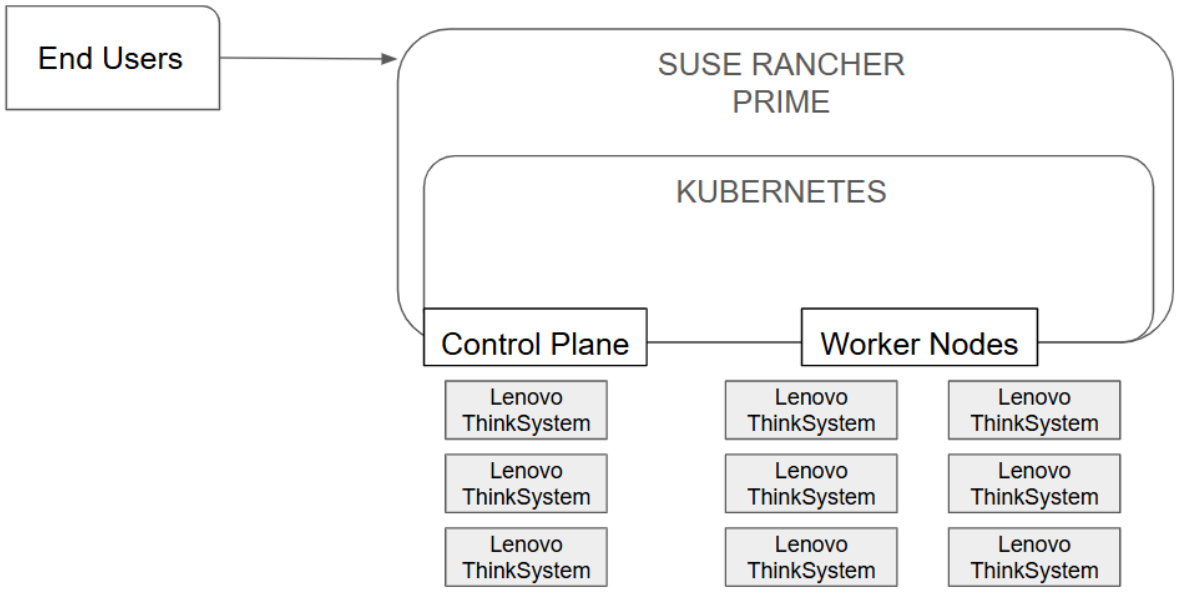

Figure 1. Architectural Overview of Kubernetes with SUSE Rancher Prime on Lenovo ThinkSystem Server Cluster

The choice of operating system is strongly related to the choice of container platform as some container platforms can only run on a small number of operating systems. See https://www.suse.com/suse-rancher/support-matrix/all-supported-versions/rancher-v2-12-1/ for more details.

A robust and flexible software stack is essential for deploying and operating enterprise-grade Kubernetes environments. This reference architecture integrates industry-standard components that provide the orchestration, management, networking, storage, and runtime capabilities needed to support modern containerized workloads. The following sections outline the key elements of the stack.

3.1 Kubernetes Orchestration

At the foundation of the stack is **Kubernetes**, the de facto standard for container orchestration. Kubernetes automates the deployment, scaling, and management of containerized applications, providing the platform consistency required to run workloads across on-premises data centers, cloud environments, and edge sites. This architecture supports both upstream Kubernetes distributions as well as lightweight variants such as **K3s**, enabling deployment in a wide range of environments with differing resource constraints.

3.2 SUSE Rancher Prime Management Platform

To simplify multi-cluster and hybrid operations, this architecture incorporates the **SUSE Rancher management platform**. SUSE Rancher Prime delivers centralized lifecycle management for Kubernetes clusters, offering:

- Unified visibility and monitoring across all clusters.
- Policy enforcement and governance for security and compliance.
- Role-based access control (RBAC) and identity management integration.
- Seamless integration with both public cloud providers and edge infrastructures.
  By abstracting cluster complexity, SUSE Rancher ensures that Kubernetes can be managed at scale with reduced operational overhead.

3.3 Supported Operating Systems

The architecture supports leading enterprise-grade Linux distributions optimized for running Kubernetes workloads, including:

- **SUSE Linux Enterprise Server (SLES)** – tightly integrated with SUSE Rancher Prime and optimized for enterprise support.
- **Red Hat Enterprise Linux (RHEL)** – widely adopted across enterprises for mission-critical applications.
- **Ubuntu** – a lightweight, flexible option favored for cloud-native workloads.
  Each distribution is fully compatible with Kubernetes and SUSE Rancher Prime, allowing organizations to select the operating system that aligns with their existing infrastructure standards, licensing models, and support agreements.

3.4 Container Runtimes

The software stack supports **OCI-compliant container runtimes**, ensuring portability and adherence to open standards. The primary supported runtime is **containerd** – the default Kubernetes runtime, optimized for performance and simplicity.

3.5 CNI Plugins

Networking within Kubernetes is enabled by the **Container Network Interface (CNI)** framework. This architecture supports a range of CNI plugins to address diverse networking requirements:

- **Canal** – combines Flannel and Calico to deliver simple, reliable networking with network policy support.
- **Calico** – offers advanced networking features including network policy enforcement, security, and scalability.
- **Cilium** – leverages eBPF for high-performance, secure, and observable networking.
  By supporting multiple options, the architecture allows IT teams to choose the networking solution that best fits workload requirements and operational priorities.

3.6 Storage Integrations

Persistent storage is a critical component for running stateful applications on Kubernetes.

To support stateful workloads in Kubernetes environments, the following persistent storage solutions are recommended:

- **SUSE Longhorn**: A lightweight, cloud-native distributed block storage system purpose-built for Kubernetes and edge deployments. Developed and maintained by SUSE, Longhorn is fully integrated with SUSE Rancher Prime and RKE2. It delivers high availability even in small cluster configurations and performs reliably on both bare-metal and virtualized infrastructure. For more information, visit https://www.longhorn.io.
- **NetApp Trident**: An open-source Container Storage Interface (CSI) driver developed by NetApp, Trident enables dynamic provisioning and management of persistent volumes for Kubernetes workloads. It integrates seamlessly with NetApp's ONTAP storage systems, which are available on AFF (All-Flash FAS), FAS (Hybrid Flash), and ASA (All-Flash SAN Array) platforms. Trident is particularly well-suited for on-premises deployments, including SAP Edge Integration Cell scenarios. For more details, visit https://www.netapp.com/trident.

4 Hardware Platforms

Selecting the right hardware platforms is critical for ensuring that Kubernetes environments deliver the required levels of performance, scalability, and efficiency. This reference architecture leverages Lenovo's industry-leading server portfolio, combining **ThinkSystem** platforms for data center deployments with **ThinkEdge** platforms for distributed edge environments. Together, these systems provide the flexibility to run containerized workloads consistently across diverse infrastructure footprints.

4.1 Lenovo ThinkSystem for Data Center Deployments

For centralized, high-density data center environments, the architecture recommends the **Lenovo ThinkSystem SR630 V4** and **SR650 V4** platforms. These servers are optimized for enterprise-grade performance, scalability, and reliability, making them ideal for running large Kubernetes clusters and mission-critical workloads.

- **ThinkSystem SR630 V4**: A 1U, dual-socket server designed for performance and density. Suitable for workloads that require high compute performance in constrained rack space.

Figure 2. Lenovo ThinkSystem SR630 V4

- **ThinkSystem SR650 V4**: A 2U, dual-socket server offering greater expandability with additional drive bays, memory, and GPU support. Well-suited for workloads that demand higher scalability, storage capacity, or heterogeneous accelerators.

Figure 3. Lenovo ThinkSystem SR650 V4

Both platforms are engineered for enterprise data centers, delivering:

- Broad CPU support for Intel® Xeon® Scalable processors and AMD EPYC™ options.
- High memory density to support in-memory workloads and large Kubernetes clusters.
- Flexible I/O configurations for diverse workload requirements, including storage- or network-intensive applications.

- Integrated management capabilities (Lenovo XClarity) for simplified monitoring and lifecycle management.

4.2 Lenovo ThinkEdge for Edge Deployments

For edge computing use cases, where space, power, and environmental conditions may be constrained, the architecture incorporates the **Lenovo ThinkEdge SE100** and **SE455 V3** platforms. These servers are specifically designed to deliver enterprise-class performance in compact and ruggedized form factors.

- **ThinkEdge SE100**: An ultra-compact edge server designed for space-constrained deployments and low-power environments. Ideal for running lightweight workloads, IoT data processing, and small-scale Kubernetes clusters with SUSE Rancher Prime and K3s at the edge.

Figure 4. Lenovo ThinkEdge SE100

- **ThinkEdge SE455 V3**: A high-performance edge server engineered for demanding workloads at scale. Optimized for AI/ML inference, real-time analytics, and large Kubernetes clusters managed by SUSE Rancher Prime, delivering enterprise-grade compute at the edge.

Figure 5. Lenovo ThinkEdge SE455 V3

Both edge platforms are designed with:

- Ruggedized chassis for operation in non-traditional IT environments (e.g., retail, manufacturing floors, telecom towers).
- Wide temperature tolerance and secure, tamper-resistant features for deployment outside secure data centers.

- Support for acceleration technologies, enabling real-time analytics and AI-driven workloads.
- Remote management and automation, allowing centralized IT teams to operate clusters consistently across thousands of distributed sites.

4.3 CPU Options and Sizing Considerations

Selecting the right processor configuration is critical for ensuring Kubernetes and Rancher management clusters perform reliably and scale effectively. Lenovo's ThinkSystem and ThinkEdge platforms offer flexible CPU options, including **Intel® Xeon® Scalable** and **AMD EPYC™** families, to meet diverse workload requirements across small, medium, and large deployments.

Based on the recommended architectures for SUSE Rancher Prime:

Tier	Target Workload	Cluster Layout	Physical Servers	Per-Node Specs	Lenovo Options
S (Small)	≤10 clusters / ≤300 nodes	3× etcd + control-plane (+2 workers optionals)	3 (minimum) +2 optional if running Rancher add-ons	CPU: 4–8 cores RAM: 16–32 GB Storage: 1× M.2 NVMe (OS/etcd), RAID1 if possible Network: 1–10 GbE	Edge: ThinkEdge SE100 DC: SR630 V4
M (Medium)	~10–50 clusters/ 300–1500 nodes	3× etcd, 2× control-plane (+2 workers optional)	5 (minimum) +2 optional for add-ons	CPU: 8–16 cores RAM: 32–64 GB Storage: NVMe SSD (≥500 IOPS), RAID1 boot Network: 10 GbE	Edge: SE455 V3 DC: SR630 V4 or SR650 V4
L (Large)	50+ clusters / >1500 nodes	3× etcd, 2× control-plane, 2× workers	7 (minimum) Increase workers if heavy add-ons	CPU: 16–32 cores RAM: 64–256 GB Storage: Dedicated NVMe for etcd + RAID1 boot Network: 25 GbE recommended	DC: SR650 V4 Edge Core: SE455 V3

Additional Considerations

- **Control Plane Nodes:** Require moderate CPU resources but prioritize stability and low latency. HA is achieved with at least three nodes hosting control plane and etcd components.
- **etcd Nodes:** The etcd datastore is latency-sensitive and critical for Kubernetes cluster health.
- **Worker Nodes:** CPU sizing depends on workload type. Compute-intensive workloads (AI/ML, analytics) benefit from higher core counts and optional GPU acceleration, while general-purpose microservices can run on mid-range configurations.
- **Edge Nodes:** For ruggedized or space-constrained environments, energy-efficient CPUs in SE100 or SE455 V3 provide flexibility for lightweight or compute-heavy edge workloads.

- **Future Growth:** Overprovisioning CPU resources is recommended to accommodate cluster expansion, workload bursts, and resource-intensive services such as monitoring or security scanning.

By aligning CPU sizing with cluster tier and workload profile, Lenovo platforms enable organizations to right-size infrastructure for performance, resiliency, and operational efficiency.

5. Networking & Storage Architecture

5.1 Onboard and SmartNIC Options

Lenovo ThinkSystem servers are equipped with onboard 10GbE and 25GbE network interface cards (NICs), providing high-throughput connectivity suitable for Kubernetes workloads. For environments requiring enhanced performance and reduced CPU load, **SmartNICs** are available as optional upgrades. These intelligent NICs can offload key networking functions such as **Container Network Interface (CNI)** processing and **Data Plane Development Kit (DPDK)** acceleration. This offloading enables more efficient packet handling, reduces latency, and frees up CPU resources for application workloads, making SmartNICs particularly valuable in high-performance, edge, or AI/ML deployments.

5.2 Network Topology and Switch Compatibility

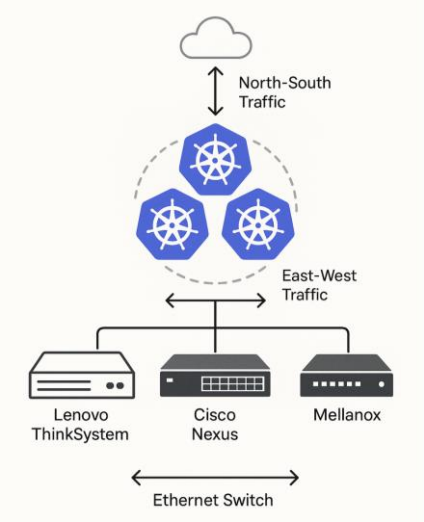

Figure 6. Network Topology and Switch Compatibility

The architecture supports **flexible and high-performance network designs**, integrating seamlessly with **Lenovo ThinkSystem Ethernet switches**, as well as **Cisco Nexus** and **Mellanox** platforms. These switches enable **high-throughput, low-latency connectivity** across Kubernetes clusters, supporting both east-west traffic within the cluster and north-south traffic for external access. This compatibility ensures that organizations can tailor their network topology to meet specific performance, scalability, and redundancy requirements, whether in a centralized data center or distributed edge environment.

5.3 Local Storage (SSDs with RAID)

Local SSDs provide high-performance, low-latency storage for Kubernetes clusters deployed on bare-metal Lenovo servers. When combined with appropriate RAID configurations, this approach delivers a balanced design that addresses both system resilience and application performance without the complexity of external storage systems.

5.3.3 Operating System Disk Layout

For each worker node, the operating system is installed on two locally attached drives configured in **RAID 1**. This configuration protects the node against single-disk failure, ensuring OS availability and simplifying node recovery while maintaining predictable performance. Using RAID 1 for the OS is a common and recommended practice in bare-metal Kubernetes environments.

5.3.4 Data Disk Layout for Kubernetes Workloads

Application data disks are presented as **Non-RAIDed (JBOD)** devices or, where required, configured with **RAID 10** for specific performance-sensitive workloads. This design is intentional when using **Longhorn**, SUSE's cloud-native distributed block storage solution.
Longhorn provides **replicated volumes across multiple Kubernetes nodes** (three replicas by default), delivering fault tolerance at the cluster level rather than relying solely on local RAID. Avoiding RAID 1 for application data disks prevents unnecessary duplication of data and ensures optimal use of storage capacity while still maintaining high availability.

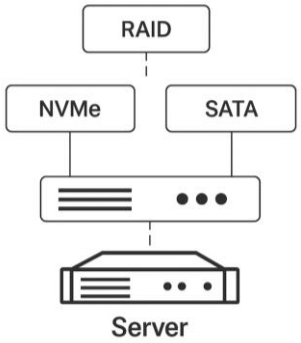

Figure 7. Local Storage (SSD with RAID)

The diagram illustrates the local storage architecture:

- **Worker nodes** equipped with SSDs.
- **RAID 1/10** setup for data protection and availability.
- **Longhorn integration** for managing persistent volumes within the cluster.
- **Storage flow** showing how data is distributed and accessed across Kubernetes workloads.

5.4 External Storage (Lenovo DE/DM/DG/DS series)

Enterprise-grade external storage is seamlessly supported through the **Trident CSI driver**, enabling Kubernetes clusters to dynamically provision and manage persistent volumes using protocols such as **iSCSI**. This integration ensures compatibility with advanced storage platforms, such as **Lenovo DE/DM/DG/DS series**, which offer centralized, scalable, and high-performance storage solutions tailored for production workloads.

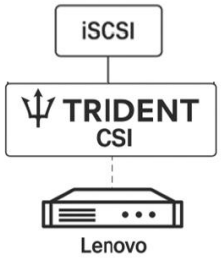

Figure 8. External Storage

These systems provide features such as deduplication, snapshots, and disaster recovery, making them ideal for mission-critical applications and hybrid infrastructure environments.

5.5 Cloud-Integrated Storage Options

SUSE Rancher Prime offers robust support for **cloud-native block storage services**, including **Amazon Elastic Block Store (EBS)**, **Microsoft Azure Disk**, and **Google Cloud Persistent Disk.**

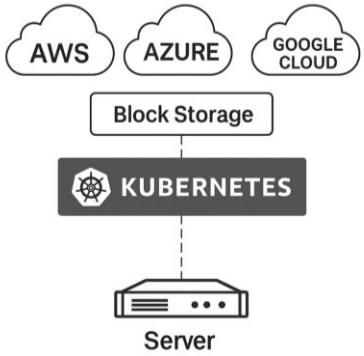

Figure 9. Cloud-Integrated Storage

These integrations enable Kubernetes clusters to dynamically provision persistent volumes in cloud environments using **CSI (Container Storage Interface) drivers**. This capability facilitates **hybrid and multi-cloud deployments**, allowing workloads to scale seamlessly across on-premises and cloud infrastructure. It also simplifies storage management by automating provisioning, resizing, and lifecycle operations, ensuring consistent performance and availability without the need to manage physical storage hardware.

6. Security & Operations

Security and operational excellence are essential when running SUSE Rancher Prime on bare-metal servers. While this deployment model provides maximum control over infrastructure, it also requires strong practices to safeguard clusters and streamline daily management. This chapter outlines how to secure Rancher-managed environments, focusing on identity and access management, auditing, compliance, and operational best practices to ensure resilience, consistency, and trustworthiness at scale.

6.1 Identity and Access Management (RBAC, External Authentication)

Identity and access management (IAM) in SUSE Rancher Prime ensures that users and applications only have the permissions necessary to perform their tasks, following the principle of least privilege. SUSE Rancher provides a multi-tenant control plane where administrators can define access policies consistently across multiple Kubernetes clusters, reducing complexity and improving governance.

Key components:

- **Role-Based Access Control (RBAC):**
  SUSE Rancher Prime integrates with Kubernetes' native RBAC to define roles and bind them to users or groups. This enables fine-grained permission control at the global, cluster, and project levels. Common patterns include creating roles for administrators, developers, and read-only users, ensuring that each role maps to clear operational responsibilities.
- **External Authentication Providers:**
  To simplify user management and align with enterprise policies, SUSE Rancher Prime supports integration with external identity providers such as Active Directory, LDAP, GitHub, and SAML-based services. This allows organizations to centralize identity lifecycle management and enforce existing corporate authentication standards (e.g., password policies, multi-factor authentication).
- **Centralized Access Policies:**
  With SUSE Rancher Prime, administrators can define global roles and consistently apply them across clusters, avoiding configuration drift. This ensures that teams have the right level of access regardless of the underlying environment, whether it's a staging cluster or production bare-metal cluster.
- **Auditing and Compliance:**
  IAM policies are a cornerstone of compliance. SUSE Rancher's integration with Kubernetes auditing capabilities enables organizations to trace user actions, verify adherence to RBAC rules, and generate evidence for regulatory requirements.

By combining Kubernetes-native RBAC with centralized identity provider integrations, SUSE Rancher Prime delivers a scalable and secure approach to identity and access management across bare-metal environments.

6.2 Cluster Hardening (CIS benchmarks, PSA policies)

Securing Kubernetes clusters is a continuous process that extends beyond initial deployment. SUSE Rancher Prime provides built-in capabilities to harden clusters according to industry best practices and compliance standards, ensuring workloads remain protected on bare-metal environments. Two key aspects of this

hardening are CIS (Center for Internet Security) benchmarks and Kubernetes Pod Security Admission (PSA) policies.

CIS Benchmarks:

- SUSE Rancher Prime includes automated scans based on CIS Kubernetes Benchmarks, enabling administrators to assess cluster configurations against standardized security guidelines.
- The scans identify misconfigurations or deviations from best practices (e.g., insecure API server flags, excessive privileges), and provide actionable remediation steps.
- Scheduled and on-demand scans help maintain continuous compliance, which is particularly important for regulated industries.

Pod Security Admission (PSA) Policies:

- PSA policies enforce security boundaries at the workload level, defining what pods can and cannot do.
- Administrators can choose predefined modes (Privileged, Baseline, Restricted) to align workloads with the desired security posture.
- For example, “Restricted” mode can prevent containers from running as root, using host networking, or mounting sensitive host paths.
- SUSE Rancher Prime enables consistent application of PSA policies across clusters, ensuring that workloads deployed on bare-metal servers adhere to organizational security standards.

SUSE Rancher Prime provides a strong foundation for cluster hardening, though CIS benchmark validation with PSA enforcement. This ensures bare-metal Kubernetes environments remain resilient, compliant, and aligned with the principle of least privilege.

6.3 Runtime Security (Trivy, NeuVector)

Securing workloads at runtime is a critical layer of defense, complementing cluster hardening and access controls. While static checks reduce risks before deployment, runtime security ensures that vulnerabilities, misconfigurations, and malicious activities are detected and mitigated as applications run on Kubernetes clusters. SUSE Rancher Prime integrates with open-source and enterprise tools such as **Trivy** and **NeuVector** to provide continuous protection.

Trivy – Vulnerability and Configuration Scanning:

- Trivy scans container images, file systems, and Infrastructure-as-Code (IaC) artifacts for known vulnerabilities, misconfigurations, and compliance issues.
- Integrated into Rancher pipelines or CI/CD workflows, it prevents insecure images from being deployed.
- Continuous scanning of running workloads ensures that new CVEs are quickly identified and addressed.

NeuVector – Runtime Threat Detection and Network Security:

- NeuVector provides deep runtime security by monitoring container and network behavior.
- It enforces **zero-trust network segmentation**, preventing unauthorized lateral movement between pods and services.

- Runtime anomaly detection identifies suspicious activities such as privilege escalations, crypto-mining, or unapproved process execution.
- NeuVector also includes compliance features like PCI DSS and HIPAA checks, which are crucial in regulated industries.

6.4 Secrets Management and image signing

Securing sensitive information and validating the integrity of software artifacts are critical to protecting Kubernetes workloads. SUSE Rancher-managed bare-metal environments should implement both **strong secrets management** and **image signing** practices.

Secrets management in Kubernetes natively allows the storage of credentials, tokens, and certificates. These are base64-encoded, offering limited protection if etcd or manifests are compromised. Because of this, the use of external secret managers such as **HashiCorp Vault, CyberArk, or cloud-native KMS** are recommended. SUSE Rancher Prime facilitates this by supporting CSI (Container Storage Interface) drivers and external secret operators.

Image signing ensures that only trusted container images are deployed to clusters. SUSE Rancher Prime supports integrations with tools like **Cosign**, which sign container images cryptographically and verify their authenticity during deployment. This mechanism prevents tampered or malicious images from being pulled into production, protecting the platform from a growing class of supply chain attacks.

The combination of strong secrets management and image signing builds trust in both directions: workloads can securely access what they need, and the platform can trust the workloads themselves. By addressing these two areas together, SUSE Rancher helps organizations on bare metal close critical security gaps that span both **data confidentiality** and **software integrity**.

6.5 Monitoring and Observability (Prometheus, Grafana, Lenovo XClarity)

Monitoring and observability are essential for operating Kubernetes clusters at scale, ensuring that both the platform and the workloads remain healthy, performant, and predictable. In a SUSE Rancher-managed bare-metal environment, combining **Kubernetes-native monitoring** with **infrastructure-level observability** provides a complete view from hardware to application.

- **Prometheus**: SUSE Rancher Prime integrates Prometheus for metrics collection across clusters, nodes, and workloads. It gathers time-series data on CPU, memory, storage, and network usage, as well as Kubernetes events and API performance.
- **Grafana**: Grafana builds on Prometheus by providing customizable dashboards and visualization. Operators and developers can track cluster performance, workload behavior, and application SLIs/SLOs in real time.
- **Alerting**: With Prometheus Alertmanager and Grafana alerts, operators can configure thresholds and notifications to proactively address issues such as resource exhaustion, failed pods, or degraded API responsiveness.

For bare-metal deployments, monitoring hardware health is as important as monitoring the Kubernetes layer. **Lenovo XClarity** provides insights into servers, power, cooling, firmware, and hardware failures. Combined with SUSE Rancher's monitoring stack enables operators to correlate cluster performance with physical server conditions, helping diagnose issues like thermal throttling, failing disks, or under provisioned hardware.

6.6 Backup and Disaster Recovery Strategies

Clusters must be prepared for failures ranging from accidental data loss to hardware breakdowns or full site outages. A robust backup and disaster recovery (DR) plan ensures business continuity and minimizes downtime.

The **backup** can be separated in three parts:

- **Etcd backups**: Regularly back up Kubernetes' key-value store (etcd), which holds all cluster state. SUSE Rancher Prime provides automated etcd snapshot management, including scheduled backups and secure remote storage.
- **Persistent volume backups**: Integrate with storage-level backup solutions or CSI snapshots to protect stateful workloads.
- **Configuration and manifests**: Store cluster and workload manifests in version control (GitOps approach), ensuring clusters can be redeployed consistently.

For disaster recovery approaches can be mentioned the following:

- **Cluster-level recovery**: Restore etcd snapshots to recover cluster state on the same or replacement nodes.
- **Application recovery**: Use persistent volume snapshots and external backup systems to restore stateful workloads.
- **Site failover**: Maintain secondary bare-metal or cloud-based clusters as cold, warm, or hot DR sites, synchronized via SUSE Rancher's multi-cluster management capabilities.
- **Testing**: Periodically validate DR plans with simulated recovery exercises to ensure procedures are effective and operators are trained.

7. Deployment Models

7.1 High-Availability Compact Cluster

The High-Availability Compact Cluster is intended to validate Kubernetes, Rancher, and associated platform components, not to provide high availability or production-grade resiliency. The focus is on simplicity, cost efficiency, and functional validation.

Key Characteristics

- Reduced hardware footprint
- Consolidated node roles
- Limited or no high availability
- Suitable for short-lived or non-critical workloads

Recommended Architecture

- **3 nodes total**, each running:
  - Kubernetes control plane components
  - Kubernetes etcd components
- **Rancher Server** deployed as a container within the cluster
- **Local storage** using:
  - Local Persistent Volumes or lightweight storage
  - Longhorn optional, with reduced replica count
- **Networking**
  - Simple CNI such as Calico or Canal
  - Standard 1 GbE connectivity is sufficient for PoC environments; bonding or redundant links are optional but not required.

Availability Considerations

- Control plane and workloads share the same nodes
- Node failure impacts both management and workloads
- Acceptable risk for evaluation and learning purposes

Use Cases

- Software validation and feature evaluation
- Rancher and Kubernetes operational learning
- Dev/Test environments
- Demonstrations and platform experimentation

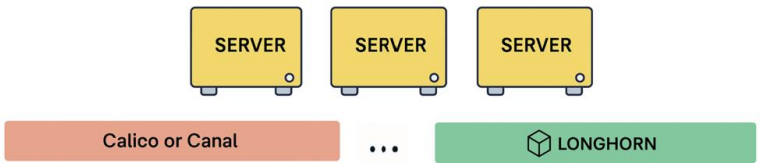

Figure 10. High-Availability Compact Cluster

7.2 High-Availability Standard Cluster

The High-Availability Standard Cluster architecture is designed for mid-sized Kubernetes environments where resilience and scalability are important, but the hardware footprint must remain balanced. This setup provides redundancy for critical components while optimizing cost and complexity compared to an enterprise-grade design.

Key Characteristics

- Redundant control plane and datastore components for fault tolerance
- Balanced hardware footprint suitable for regional or departmental deployments
- Scalable to support moderate workloads and multiple downstream clusters
- Appropriate for production environments with non-critical or moderately critical workloads

Recommended Architecture

- **Minimum 5 nodes total**, including:
  - **3 nodes** dedicated to Kubernetes control plane and etcd components
  - **2 nodes** dedicated to Kubernetes worker roles for Rancher add-ons and workloads
- **External Load Balancer** (e.g., HAProxy or NGINX) for API traffic distribution and failover
- **Storage**:
  - Local SSDs for etcd and OS
  - Optional lightweight storage solution such as Longhorn with standard replica count
- **Networking**:
  - 10GbE recommended for inter-node communication
  - Redundant NICs or bonded interfaces for improved availability
- **Power**:
  - Dual power supplies per server connected to separate circuits
- **Management**:
  - BMC/IPMI interfaces for out-of-band management and recovery

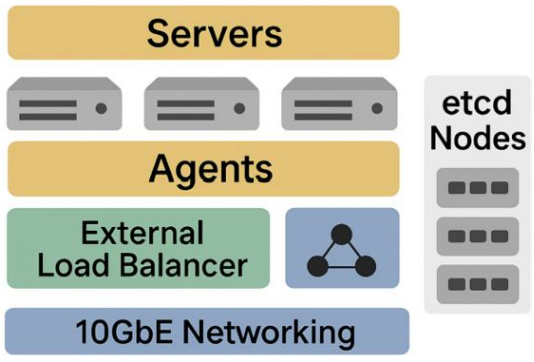

Figure 11. High-Availability Standard Cluster

Availability Considerations

- Redundant control plane and datastore components prevent single points of failure
- Supports maintenance with minimal downtime
- Provides sufficient capacity for Rancher management and moderate workloads

Use Cases

- Regional or departmental Kubernetes environments under SUSE Rancher
- Production workloads requiring moderate resilience
- Environments with multiple downstream clusters and Rancher add-ons (e.g., monitoring, logging)
- Balanced deployments where cost and availability must be optimized

7.3 High-Availability Enterprise Cluster

The **High Availability (HA) Enterprise Cluster** architecture is designed for large-scale, mission-critical Kubernetes environments where uptime, performance, and resilience are essential. Unlike a minimal cluster, this design introduces **redundancy** across every major layer: control plane, storage, networking, and power, to ensure uninterrupted operation even during failures or maintenance.

Key Characteristics

- Full redundancy across control plane, storage, networking, and power
- Designed for high uptime and fault tolerance
- Scalable for large workloads and enterprise environments
- Suitable for production-grade Kubernetes deployments

Recommended Architecture

- **Minimum 7 nodes total**, including:
    - At least **2 control plane nodes** for Kubernetes control plane components
    - At least **3 dedicated etcd nodes** for isolating the Kubernetes datastore
    - At least **2 worker nodes** for Kubernetes workloads
- **External Load Balancer** (e.g., HAProxy or F5) to distribute API traffic and maintain accessibility during node failures
- **Enterprise Storage** using:
    - SAN-based solutions or systems like Lenovo DE/DG/DM/DS series integrated via Trident
    - High-performance, redundant, snapshot-capable persistent storage
- **Networking**:
    - High-speed links (25GbE or faster) with redundant paths
- **Power**:
    - Dual power supplies per server connected to separate circuits
- **Management**:

- BMC/IPMI interfaces for out-of-band control and rapid recovery

Availability Considerations

- Redundant components across all layers prevent single points of failure
- Supports maintenance without downtime
- Improves scalability and responsiveness for growing workloads

Use Cases

- Enterprise-grade Kubernetes environments under SUSE Rancher
- Continuous operation for mission-critical applications
- Simplified upgrades and lifecycle management
- Large-scale deployments requiring high resilience and performance

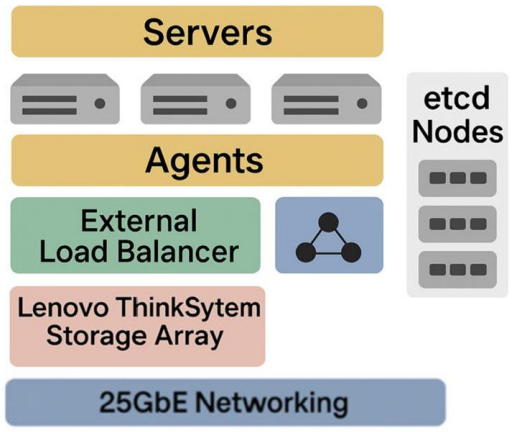

Figure 12. High-Availability Enterprise Cluster

The HA Enterprise Cluster architecture enables:

- **Seamless Upgrades and Maintenance:** Nodes or components can be serviced without affecting workloads, thanks to redundancy across layers.
- **Scalable Growth:** Adding new worker nodes or expanding storage capacity can be performed dynamically as workloads evolve.
- **Enhanced Security and Observability:** Dedicated infrastructure supports hardened security configurations, centralized logging, and detailed monitoring of cluster health.
- **Disaster Recovery Readiness:** With properly configured replication and external backups, this architecture forms the foundation for DR strategies across multiple data centers.

7.4 Edge Cluster Seployment with K3s

In modern infrastructure design, workloads are no longer confined to centralized data centers. Increasingly, enterprises require compute capabilities at the *edge*. The **Edge Cluster Deployment with K3s** architecture is

built precisely for this purpose: it delivers a lightweight, secure, and manageable Kubernetes platform optimized for resource-constrained or distributed environments.

In an edge deployment, the goal is to run applications closer to where data is generated — reducing latency, improving reliability in case of intermittent network connectivity, and ensuring faster local decision-making. SUSE Rancher Prime plays a crucial role here as the **centralized management plane**, allowing organizations to deploy, monitor, and maintain hundreds or even thousands of remote K3s clusters from one unified interface.

A typical K3s edge cluster consists of a small number of nodes — often **one to three servers** per site. The cluster can operate in two main modes:

- **Single-Node Cluster:**
  Ideal for lightweight workloads or environments with limited hardware. The server node hosts both the control plane and worker components.
- **HA (High Availability) Multi-Node Cluster:**
  Recommended for critical edge use cases. It typically includes **three K3s server nodes** (for control plane HA) and one or more **agent nodes** (workers). The datastore can be either an embedded etcd (for small deployments) or an **external datastore** such as MySQL, PostgreSQL, or etcd cluster shared among servers.

K3s clusters usually run on **bare-metal, virtual machines, or ARM-based edge devices**, and can even operate in disconnected or intermittently connected modes, syncing back to the central Rancher management plane when the network becomes available.

8. Scalability & Lifecycle Management

As Kubernetes environments grow from small proof-of-concept clusters to globally distributed enterprise deployments, the ability to **scale efficiently and manage the entire lifecycle** of clusters becomes essential. In large infrastructures managed by SUSE Rancher Prime, scalability is not just about adding more nodes — it's about ensuring consistent performance, streamlined updates, and simplified operations across diverse environments.

8.1 Node Scaling and Workload Rescheduling

In dynamic Kubernetes environments, workloads and infrastructure must continuously adapt to changing demand. **Node scaling** and **workload rescheduling** are key mechanisms that ensure the cluster remains efficient, resilient, and cost-effective as usage patterns evolve.

Node scaling refers to the process of **adding or removing** compute nodes within a cluster based on resource utilization or workload requirements. In a SUSE Rancher–managed environment, scaling can be performed **manually** or **automatically**. SUSE Rancher Prime provides **an intuitive interface** to increase or decrease node counts, and it integrates with underlying infrastructure providers, whether cloud, virtualized, or bare-metal, to provision new nodes or decommission unused ones seamlessly. When **auto-scaling** is enabled, the system can automatically adjust capacity by monitoring CPU, memory, or custom metrics to maintain optimal performance without overprovisioning resources.

When new nodes are added to the cluster, Kubernetes schedules **newly created pods** onto the additional capacity based on available resources, node conditions, taints and tolerations, affinity rules, and declared resource requests. Existing running workloads are **not automatically rebalanced or migrated** to newly added nodes.

If a node fails or is intentionally drained for maintenance, the scheduler automatically reschedules affected pods onto healthy nodes that meet the workload's placement constraints. This behavior minimizes service disruption and helps maintain application availability without requiring manual intervention.

For environments where proactive workload redistribution is required (for example, to better utilize newly added capacity), administrators must rely on controlled actions such as rolling restarts, descaling and rescaling deployments, or external tools and policies designed for workload rebalancing.

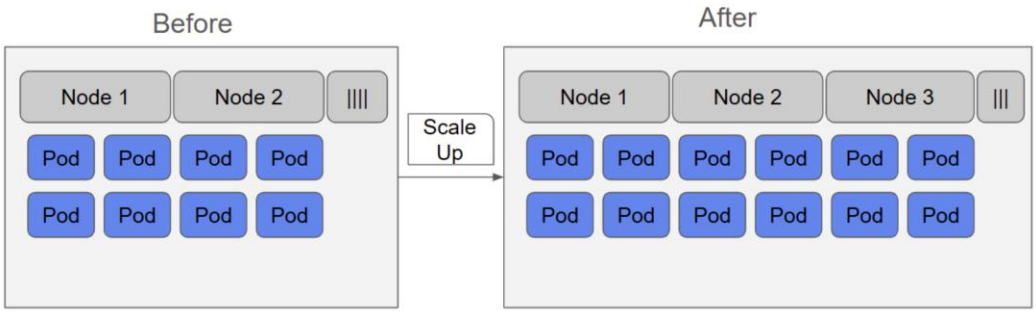

Figure 13. Workload Rescheduling

In edge and hybrid deployments, SUSE Rancher Prime extends these capabilities across multiple clusters, ensuring consistent scaling and rescheduling policies. Whether running in a centralized data center, at the edge, or in the cloud, workloads are intelligently managed to match available capacity and operational priorities. The result is a **self-healing**, adaptive environment where resources flex dynamically to meet business demands while maintaining reliability and efficiency.

8.2 Cluster Upgrades and Maintenance

Maintaining a Kubernetes environment in production requires a careful and structured approach to upgrades and ongoing maintenance. As clusters evolve over time, keeping them up to date ensures **security, stability, and compatibility** with modern workloads and integrations. SUSE Rancher Prime simplifies this process through centralized, automated lifecycle management, making upgrades safer and more predictable across multiple clusters.

Cluster upgrades typically involve updating **Kubernetes versions, system components, and node operating systems**. In a Rancher-managed environment, administrators can trigger upgrades directly from the Rancher interface or via automation pipelines, ensuring consistent and validated processes across all managed clusters. SUSE Rancher orchestrates upgrades in a **rolling fashion**, updating control plane components and worker nodes sequentially to avoid downtime. Workloads are gracefully drained and rescheduled during the process, maintaining application availability while nodes are updated.

Maintenance also includes the regular application of **security patches, configuration updates, and resource optimization**. Monitoring tools integrated with SUSE Rancher provide real-time health checks, resource usage insights, and alerts for proactive maintenance.

For enterprise and edge deployments, automated scheduling of maintenance windows and version policies ensures clusters remain compliant and up to date without impacting critical services. By managing upgrades, patching, and configuration consistency at scale, SUSE Rancher Prime transforms what would traditionally be a complex, high-risk operation into a streamlined, auditable, and repeatable process.

In essence, cluster upgrades and maintenance under Rancher are designed to keep environments current and resilient — ensuring continuous delivery of applications while minimizing operational overhead and human error.

8.3 Integration with CI/CD and DevSecOps Pipelines

To deliver applications quickly and safely, modern IT teams rely on **automation** — moving code from development to production through continuous integration and delivery (CI/CD). When combined with security practices (DevSecOps), this process ensures that every update is tested, verified, and deployed in a secure and repeatable way.

SUSE Rancher Prime makes it easy to connect Kubernetes clusters to existing CI/CD tools such as **Jenkins, GitLab CI, Argo CD, Tekton, or GitHub Actions**. This allows new applications or updates to be automatically deployed into Rancher-managed clusters as soon as they pass tests or are approved in version control systems. Instead of manually managing configurations, teams can use a **GitOps approach**, where everything from cluster setup to application settings, is stored in Git and automatically applied when changes are made.

9. Licensing & Support

9.1 SUSE Rancher Prime

SUSE Rancher Prime is a commercially supported edition that includes enterprise-grade features, long-term support, and access to SUSE subscription services. Prime offers enhanced security, certified integrations, and professional support from SUSE, ensuring higher reliability and faster resolution of issues in production environments.

The key benefits of SUSE Rancher Prime are:

- **Enterprise Support:** 24/7 SLA-backed support, bug fixes, upgrade validation, knowledge base access.
- **Security & Compliance:** Advanced security, automated cluster hardening, FIPS 140-2 and other compliance certifications.
- **Extended Lifecycle:** Long-term support for RKE2/K3s, access to premium features, regular updates.
- **Observability & Monitoring:** Real-time insights, easier troubleshooting, integration with monitoring tools.

9.2 SUSE Support Options

SUSE offers a comprehensive support framework for SUSE Rancher Prime, designed to ensure smooth, reliable, and secure operations of your Kubernetes environment. At the core, there are different support models tailored to organizational needs.

The **Standard Support** model provides essential assistance during business hours, helping teams quickly address common issues and maintain operational continuity. For organizations requiring uninterrupted access, **Priority Support** offers 24/7 availability, guaranteeing rapid response times and direct engagement with SUSE experts.

Beyond availability, SUSE support encompasses a wide range of services to maintain system health and performance. This includes **root cause analysis** for diagnosing complex problems, **upgrade validation** to ensure safe and smooth updates, and periodic **Rancher supportability reviews** to optimize your deployment. Customers also gain access to a **private Slack channel** for direct communication with SUSE engineers, a **trusted image registry** for secure container images, and an extensive **knowledge base** filled with best practices and guidance.

To further strengthen the solution, the SUSE Rancher Prime Suite offering includes support for all the products in the suite, including SUSE Virt., SUSE Storage, SUSE Security, SUSE Observability, and more. This integrated support ecosystem ensures that every component of your Rancher Prime deployment is covered, providing both peace of mind and enterprise-grade reliability.

Appendix: Lenovo Bill of Materials

This appendix contains the bill of materials (BOMs) for computational servers.

ThinkEdge SE100

Part number	Product Description	Qty
7DGRCTO1WW	Node : ThinkEdge SE100 - 3 Year Base Warranty	1
C31D	ThinkEdge SE100 Chassis	1
C31W	ThinkEdge SE100 Planar with Intel Core Ultra 7 255H ,16C, 28W, 2.0GHz	1
C39J	ThinkEdge 16GB TruDDR5 6400MHz CSODIMM	1
C8V2	ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	1
AUKP	ThinkSystem Broadcom 57416 10GBASE-T 2-Port PCIe Ethernet Adapter	1
C39R	ThinkEdge 140W 230V/115V External Power Supply	1
6313	2.8m, 10A/120V, C13 to NEMA 5-15P (US) Line Cord	1
C31C	ThinkEdge SE100 Port Dust Cover Kit	1
B755	Desktop Mode	1
C31J	ThinkEdge SE100 Bottom Rubber Feet	1
C31A	ThinkEdge SE100 Fan Module	1
BYBQ	XClarity Controller Managed	1
B0MK	Enable TPM 2.0	1
B7XZ	Disable IPMI-over-LAN	1
BB98	Disable IPMI-over-KCS	1
A2N7	Planar Not Integrated With Chassis	1
B0ML	Feature Enable TPM on MB	1
BRPJ	XCC Platinum	1
C3GN	ThinkEdge SE100 LPK	1
BE0B	Non-Redundant	1
C305	ThinkEdge SE100 REGID	1
C8U9	Top-Cover Thermal Gap Pad Kit	1
C8UA	Bottom-Cover Thermal Gap Pad Kit	1
C308	ThinkEdge SE100 M.2 Holder	1
C8UB	Expansion Kit Rubber Feet	1
C316	ThinkEdge SE100 Fan Cable (Bridge Cable)	2
C31P	ThinkEdge SE100 Node SSL_LI	1
C31M	ThinkEdge SE100 Node Label GBM	1
C31K	ThinkEdge SE100 Node WW Packaging	1
C31T	ThinkEdge SE100 Agency Label	1
C8UC	Front I/O Panel	1
C319	ThinkEdge SE100 Node Cosmetic Cover	1
C30E	SE100 Expansion Kit for NIC Adapter	1
7S0XCTO5WW	XClarity Controller Platin-FOD	1
SBCV	Lenovo XClarity XCC2 Platinum Upgrade (FOD)	1

ThinkEdge SE455 V4

Part number	Product Description	Qty
7DBYCTO1WW	Server : ThinkEdge SE455 V3-3yr Base Warranty	1
BVTK	ThinkEdge SE455 V3 Chassis	1
BY92	Optimised for Acoustics - Mode 1	1
BY8S	System Operational Temperature 5C to 25C / 41F to 77F	1
BW2V	ThinkEdge SE455 V3 AMD EPYC 8124P 16C 125W 2.45GHz Processor	1
BW3M	ThinkEdge SE455 V3 2U Heatsink	1
BQ39	ThinkSystem 32GB TruDDR5 4800MHz (1Rx4) 10x4 RDIMM-A	2
5977	Select Storage devices - no configured RAID required	1
BMFT	ThinkSystem RAID 540-8i PCIe Gen4 12Gb Adapter	1
CABQ	ThinkSystem 2.5" VA 3.2TB Mixed Use SAS 24Gb HS SSD	3
BYLR	ThinkSystem 2.5" VA 480GB Read Intensive SATA 6Gb HS SSD v2	2
BVVN	ThinkEdge SE455 V3 Internal 2.5" Drive Cage	1
BVUV	ThinkEdge SE455 V3 4x2.5" SAS/SATA Backplane	1
BVUV	ThinkEdge SE455 V3 4x2.5" SAS/SATA Backplane	1
B5T1	ThinkSystem Broadcom 5719 1GbE RJ45 4-port OCP Ethernet Adapter	1
BK1H	ThinkSystem Broadcom 57414 10/25GbE SFP28 2-port PCIe Ethernet Adapter	1
BVUR	ThinkEdge SE455 V3 Riser1	1
BMH8	ThinkEdge 1100W 230V/115V Platinum Hot-Swap Gen2 Power Supply	1
BMH2	ThinkEdge 600mm Ball Bearing Rail Kit	1
BS4E	ThinkEdge 130mm USB-C to VGA Display Cable	1
BYBQ	XClarity Controller Managed	1
BVV6	ThinkEdge SE455 V3 Intrusion Switch	1
BVTX	ThinkEdge SE455 V3 Standard EIA Bracket	1
BVTL	ThinkEdge SE455 V3 Motherboard	1
BRPJ	XCC Platinum	1
BK14	Low voltage (100V+)	1
BVUA	ThinkEdge SE455 V3 Language Pack	1
BVV4	ThinkEdge SE455 V3 1st Motherboard Power Cable	1
BVV5	ThinkEdge SE455 V3 2nd Motherboard Power Cable	1
BVV2	ThinkEdge SE455 V3 Fan Board Sideband Cable	1
BVV3	ThinkEdge SE455 V3 Fan Board Power Cable	1
BVVW	ThinkEdge SE455 V3 Front Backplane Power Cable	1
BVUX	ThinkEdge SE455 V3 Internal Backplane Power Cable	1
BVVV	ThinkEdge SE455 V3 X40 SAS/SATA Y-Cable for Front and Internal Drive Bays	1
BW1P	ThinkEdge SE455 V3 Riser1 PCIe Gen5 Cable	1
BW1Q	ThinkEdge SE455 V3 Riser1 PCIe Gen4 Cable	1
C8YD	SE455 V3 Laser Service Indicator	1
BVTP	ThinkEdge SE455 V3 Fan	5
BVUH	ThinkEdge SE455 V3 Dust Cover Kit for I/O Ports	1

BVUJ	ThinkEdge SE455 V3 Regulatory ID	1
BVUK	ThinkEdge SE455 V3 Power Distribution Board	1
BVUL	ThinkEdge SE455 V3 Power Distribution Board Sideband Cable	1
BVV1	ThinkEdge SE455 V3 Fan Board	1
BW38	ThinkEdge SE455 V3 Supercap Holder	1
BW36	ThinkEdge SE455 V3 M.2 Air Baffle	1
BW3A	ThinkEdge SE455 V3 CPU Air Baffle for 2U Heatsink	1
BVUT	ThinkEdge SE455 V3 Riser2 Filler	1
BVVF	ThinkEdge SE455 V3 Riser Side Support	3
BVUN	ThinkEdge SE455 V3 Power Supply Filler	1
BW3E	ThinkEdge SE455 V3 1100W Platinum PSU Rating Label	1
BVU8	ThinkEdge SE455 V3 LI Service Label	1
BVU4	ThinkEdge SE455 V3 Label Group	1
BVUP	ThinkEdge SE455 V3 PSU Agency Label : CCC, CECP	1
BW3N	ThinkEdge SE455 V3 Security Activation Label	1
BVTM	ThinkEdge SE455 V3 Root of Trust	1
BVU3	ThinkEdge SE455 V3 System Package	1
7S0XCTO5WW	XClarity Controller Platin-FOD	1
SBCV	Lenovo XClarity XCC2 Platinum Upgrade (FOD)	1
5641PX3	XClarity Pro, Per Endpoint w/3 Yr SW S&S	1
1340	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1

ThinkSystem SR630 V4

Part number	Product Description	Qty
7DG9CTO1WW	Server : ThinkSystem SR630 V4-3yr Base Warranty	1
C1XE	ThinkSystem 1U V4 10x2.5" Chassis	1
C3JB	ThinkSystem General Computing - Power Efficiency	1
BVGL	Data Center Environment 30 Degree Celsius / 86 Degree Fahrenheit	1
C5QR	Intel Xeon 6520P 24C 210W 2.4GHz Processor	2
C1XJ	ThinkSystem 1U V4 Performance Heatsink	2
C0TQ	ThinkSystem 64GB TruDDR5 6400MHz (2Rx4) RDIMM	2
5977	Select Storage devices - no configured RAID required	1
C0TU	ThinkSystem RAID 545-8i PCIe Gen4 12Gb Adapter	1
BYLR	ThinkSystem 2.5" VA 480GB Read Intensive SATA 6Gb HS SSD v2	2
BYLT	ThinkSystem 2.5" VA 1.92TB Read Intensive SATA 6Gb HS SSD v2	3
C21T	ThinkSystem 1U V4 8x2.5" SAS/SATA Backplane	1
C1YK	ThinkSystem SR650 V4/SR630 V4 x16 OCP Cable Kit	1
B5T1	ThinkSystem Broadcom 5719 1GbE RJ45 4-port OCP Ethernet Adapter	1
BN2T	ThinkSystem Broadcom 57414 10/25GbE SFP28 2-Port OCP Ethernet Adapter	1
C1YH	ThinkSystem SR630 V4 x16/x16 PCIe Gen5 Cable Riser 1	1

C1Z7	ThinkSystem SR630 V4 Full Height+Low Profile Riser1 Cage	1
C0U5	ThinkSystem 1300W 230V/115V Platinum CRPS Hot-Swap Power Supply v2.4	2
6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord	2
C1YT	ThinkSystem 1U V4 Performance Fan Module	4
C1YP	ThinkSystem 1U V4 Standard Media Bay	1
C2DH	ThinkSystem Toolless Slide Rail Kit V4	1
BK14	Low voltage (100V+)	1
B8KV	ThinkSystem 1U 8x2.5" SAS/SATA HDD Type Label	1
BZ7F	ThinkSystem WW Lenovo LPK, Birch Stream	1
B97B	XCC Label	1
C212	ThinkSystem BHS SR630 V4 1U PCIe number 1-2 and OCP number Label (BF+Rear M.2)	1
C1ZS	ThinkSystem 1300W (CRPS) power rating label WW	1
AWF9	ThinkSystem Response time Service Label LI	1
C20D	ThinkSystem SR630 V4 model name Label	1
C20J	ThinkSystem SR630 V4 Service Label for WW	1
AUTQ	ThinkSystem small Lenovo Label for 24x2.5"/12x3.5"/10x2.5"	1
BQPS	ThinkSystem logo Label	1
C1ZP	ThinkSystem SR630 V4 Agency label with Blank	1
C1XZ	ThinkSystem Power Cable, PIC Power 2x3+6P to PIC Power 2x3+6P, 380mm	1
C1XV	ThinkSystem G4 SAS Cable, SLX8-SLX4 x2, 900mm/1020mm	1
BE0F	N+N Redundancy Without Over-Subscription	1
CAR5	SR630 V4 Laser service indicator	1
CA9B	ThinkSystem 1U/2U V4 From MB-PCIe10 to MB OCP1 EXPANSION	1
C3NP	ThinkSystem SR630 V4 MI-BF Cbl Riser to P9	1
B0ML	Feature Enable TPM on MB	1
C4DV	ThinkSystem SR630 V4 MotherBoard	1
C3K9	XClarity Platinum Upgrade v3	1
BU1E	Top Choice flag	1
CA7N	ThinkSystem SR630 V4 System I/O board v2	1
C26Y	ThinkSystem V4 CPU HS Clip	2
BHSS	MI for PXE with RJ45 Network port	1
C5XG	ThinkSystem SR630 V4 General Config PKG AC+CL	1
B5X6	ThinkSystem 1U 2x2.5" Fixed Filler	1
AVEN	ThinkSystem 1x1 2.5" HDD Filler	3
B8NK	ThinkSystem 1U Super Cap Holder Dummy	1
C1YY	ThinkSystem 1U V4 Low Profile Riser Cage Filler	1
C2ZB	ThinkSystem RAID 545-8i PCIe Gen4 12Gb Adapter Placement	1
7S0XCTO8WW	XClarity Controller Prem-FOD	1
SCY0	Lenovo XClarity XCC3 premier - FOD	1
5641PX3	XClarity Pro, Per Endpoint w/3 Yr SW S&S	1

ThinkSystem SR650 V4

Part number	Product Description	Qty
7DGDCTO1WW	Server : ThinkSystem SR650 V4-3yr Base Warranty	1
C3QK	ThinkSystem SR650 V4 24x2.5" Chassis	1
C3JB	ThinkSystem General Computing - Power Efficiency	1
BVGL	Data Center Environment 30 Degree Celsius / 86 Degree Fahrenheit	1
C5QT	Intel Xeon 6530P 32C 225W 2.3GHz Processor	2
BPDR	ThinkSystem V4 2U Standard Heatsink	2
C0U9	ThinkSystem 32GB TruDDR5 6400MHz (1Rx4) RDIMM	8
5977	Select Storage devices - no configured RAID required	1
C0TU	ThinkSystem RAID 545-8i PCIe Gen4 12Gb Adapter	1
BYLR	ThinkSystem 2.5" VA 480GB Read Intensive SATA 6Gb HS SSD v2	2
BYM5	ThinkSystem 2.5" VA 1.92TB Mixed Use SATA 6Gb HS SSD v2	6
C3RU	ThinkSystem 2U V4 8x2.5" AnyBay Backplane	1
B5T1	ThinkSystem Broadcom 5719 1GbE RJ45 4-port OCP Ethernet Adapter	1
BN2T	ThinkSystem Broadcom 57414 10/25GbE SFP28 2-Port OCP Ethernet Adapter	1
C62D	ThinkSystem SR650/a V4 x16 Rear Direct Riser Slot 5	1
C0U5	ThinkSystem 1300W 230V/115V Platinum CRPS Hot-Swap Power Supply v2.4	2
6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord	2
C3RQ	ThinkSystem 2U 6038 24K Standard Fan Module	6
C2DH	ThinkSystem Toolless Slide Rail Kit V4	1
BQQ2	ThinkSystem 2U V3 EIA Latch Standard	1
BPKR	TPM 2.0	1
B7XZ	Disable IPMI-over-LAN	1
C3K9	XClarity Platinum Upgrade v3	1
C4S2	ThinkSystem SR650 V4 Processor board,BHS,DDR5,Santorini,2U	1
CB2P	SR650 V4 Laser service indicator	1
B0ML	Feature Enable TPM on MB	1
AVEQ	ThinkSystem 8x1 2.5" HDD Filler	2
AURS	Lenovo ThinkSystem Memory Dummy	24
BU1E	Top Choice flag	1
BHSS	MI for PXE with RJ45 Network port	1
C3RN	ThinkSystem 2U Main Air Duct	1
C26Y	ThinkSystem V4 CPU HS Clip	2
C3S5	ThinkSystem 2U V4 3FH Riser Cage	1
C4SH	HV 2U V4 General WW L1 PKG BOM	1
C3RJ	ThinkSystem 2U 2LP Riser Cage Filler	2
C7Y8	ThinkSystem SR650 V4 System I/O Board	1
C3RH	ThinkSystem 2U 3FH Riser Cage Filler	1
C3R0	ThinkSystem Power Cable, 2x6+12 P-2x3+6 P, 250 mm	1
C6TH	Think System,PCIe Gen5 Cable, MCIOx8-MCIOx8, 350 mm	4

C3QX	ThinkSystem G4 SAS Cable, SLX8-SLX8, 900mm	1
C3T9	ThinkSystem SR650 V4 model name Label	1
C1ZS	ThinkSystem 1300W (CRPS) power rating label WW	1
C3SR	ThinkSystem SR650 V4 Agency label With Blank	1
AWF9	ThinkSystem Response time Service Label LI	1
C3TH	ThinkSystem SR650 V4 Service Label for WW	1
B97B	XCC Label	1
B8K5	ThinkSystem 2U MS 24x2.5" AnyBay HDD Type Label1	1
AUTQ	ThinkSystem small Lenovo Label for 24x2.5"/12x3.5"/10x2.5"	1
BQPS	ThinkSystem logo Label	1
BZ7F	ThinkSystem WW Lenovo LPK, Birch Stream	1
CBDC	ENERGY STAR Certification Country	1
BE0F	N+N Redundancy Without Over-Subscription	1
BK14	Low voltage (100V+)	1
5374CM1	Configuration Instruction	1
C2ZB	ThinkSystem RAID 545-8i PCIe Gen4 12Gb Adapter Placement	1
7S0XCTO8WW	XClarity Controller Prem-FOD	1
SCY0	Lenovo XClarity XCC3 premier - FOD	1
5641PX3	XClarity Pro, Per Endpoint w/3 Yr SW S&S	1
1340	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1

Resources

Lenovo ThinkEdge SE100 Edge Server Product Guide

Lenovo ThinkEdge SE455 V3 Server Product Guide

Lenovo ThinkSystem SR630 V4 Server Product Guide

Lenovo ThinkSystem SR650 V4 Server Product Guide

Lenovo ThinkSystem DE Series Storage Datasheet

Lenovo ThinkSystem DG Series Storage Datasheet

Lenovo ThinkSystem DM Series Storage Datasheet

Lenovo ThinkSystem DS Series Storage Datasheet

Document History

Version 1.0	13 January 2026	Initial version

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®
ThinkEdge®
ThinkSystem®
XClarity®

The following terms are trademarks of other companies:

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

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

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

HashiCorp® is a trademark of IBM in the United States, other countries, or both.

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