Reference Architecture for SAP Edge Integration Cell on ThinkSystem V4 Servers

Last update: **February 2026**
Version: **1.0**

Provides a technical overview of SAP EIC solution

Contains sizing recommendations

Illustrates SAP EIC implementation in a production setting

Configure your system for best use with SAP EIC

Silviu Eduard Lungu
Vlad Marian Spoiala

LENOVO PRESS

Table of Contents

1 Introduction

The intended audience for this document includes enterprise architects, IT decision-makers and SAP system administrators who are evaluating or planning to deploy SAP Edge Integration Cell (EIC) on Lenovo ThinkSystem servers. This reference architecture provides guidance for building a reliable, scalable, and high-performance platform to support hybrid integration scenarios where local execution of integration content is required.

Enterprises today face an accelerating pace of digital transformation, driven by an ever-growing number of applications, devices, and data sources. While cloud-first strategies deliver scalability and innovation, many organizations still rely on on-premises systems to meet regulatory, compliance, latency, or business continuity requirements. To bridge these environments, a hybrid integration approach is essential—one that seamlessly connects cloud, on-premises, and edge systems. This paper uses the term Edge in the way SAP uses it, which is everything outside of SAP Cloud. This includes both customer on-premise environments as well as public cloud environments like AWS, Azure, or Google Cloud.

SAP Integration Suite provides this foundation through a comprehensive, enterprise-grade integration platform as a service (iPaaS), enabling organizations to unify applications, data, and processes across hybrid and multi-cloud landscapes with agility, security, and scalability.

SAP Edge Integration Cell extends these capabilities with a secure, containerized runtime for executing integration flows locally, supporting continuous operations in low-connectivity environments, ensuring data residency compliance, and enabling near real-time processing close to business operations

Lenovo ThinkSystem servers provide the optimal infrastructure foundation for SAP EIC deployments. With industry-leading performance, energy efficiency, and proven reliability, ThinkSystem servers ensure consistent execution of integration workloads while simplifying operations and lifecycle management. Designed for enterprise-grade resiliency, Lenovo servers deliver a secure, efficient, and flexible platform for hybrid integration landscapes.

This document provides architectural guidance and best practices for deploying SAP EIC on Lenovo ThinkSystem servers. It highlights the benefits of Lenovo infrastructure for hybrid integration use cases, including:

- High performance and scalability to meet the demands of containerized integration runtimes.
- Resilience and business continuity through enterprise-class reliability and availability.
- Secure and resilient server platforms designed to help organizations meet regulatory and data residency obligations.
- Flexible deployment models enabling consistent operation across edge and on-premises environments.
- Global services and support to ensure simplified implementation and long-term operational stability.

The reference architecture presents design considerations, operational models, and example configurations to guide clients in successfully deploying SAP EIC on Lenovo ThinkSystem servers. By following these recommendations, enterprises can establish a future-ready hybrid integration platform that combines the innovation of SAP’s cloud services with Lenovo ThinkSystem infrastructure.

2 Business problem

SAP Integration Suite is a single, comprehensive, and scalable Integration Platform as a Service (iPaaS) solution from SAP. It is a key part of the SAP Business Technology Platform (SAP BTP) and represents SAP's strategic, cloud-based approach to connecting business systems. Its main function is to unify, automate, and secure an organization's entire business ecosystem by handling every integration requirement in one place.

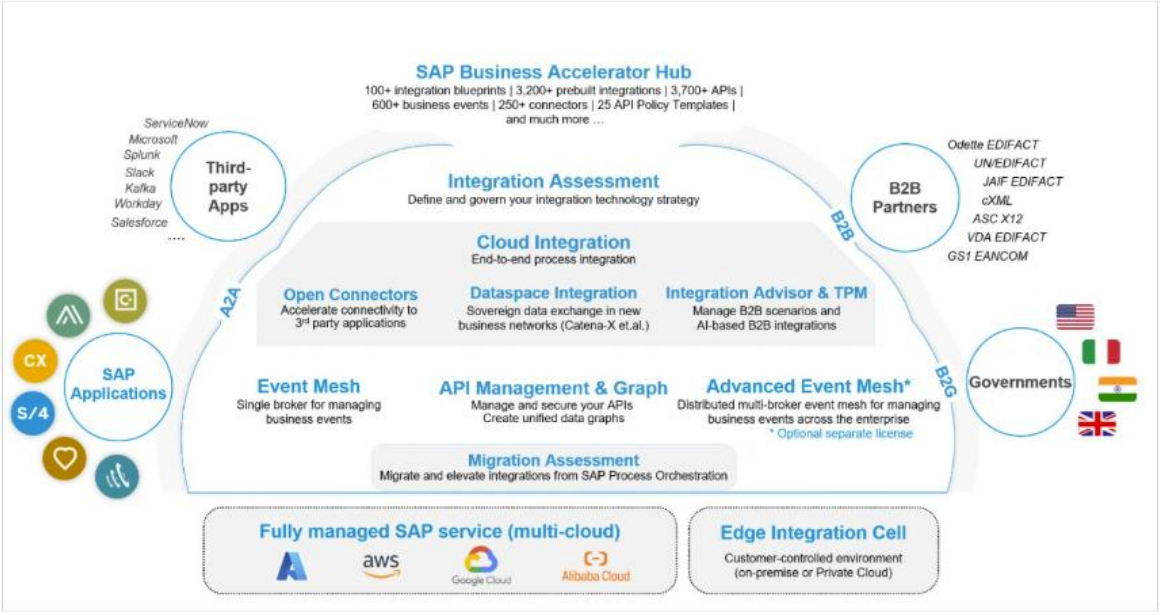

Figure 1: SAP Integration Suite and its components

SAP Integration Suite is the successor to SAP's previous on-premise integration solutions, SAP Process Integration (PI) and SAP Process Orchestration (PO). While SAP PI/PO is on-premise solution mainly focused on traditional A2A (Application-to-Application) integration, SAP Integration Suite is a cloud-native iPaaS designed for modern hybrid and multi-cloud environments, focusing on API-first, event-driven, and business-to-business integration.

An on-premise software deployment, hosted on a customer's own servers, offers maximum control over the entire IT stack, allowing for deep customization, system maintenance scheduling and data residency compliance, making it ideal for large enterprises with complex, highly specialized, or heavily regulated business processes. Conversely, the cloud approach shifts the burden of infrastructure, maintenance, and mandatory upgrades entirely to the cloud. While cloud subscriptions offer operational efficiency, they involve less control and higher latency.

SAP Edge Integration Cell (EIC) is designed to solve this challenge of hybrid integration, which is the need to combine the agility and innovation of the cloud with the security and low-latency performance of on-premise execution. This is essential for organizations that must comply with strict data sovereignty and data residency laws, or for mission-critical processes where data cannot leave the corporate firewall. SAP EIC serves as a modern, containerized runtime for the cloud-based SAP Integration Suite, allowing businesses to design and govern all their integration flows centrally in the cloud, while deploying and executing sensitive, high-volume scenarios securely and locally within their own datacenter.

Lenovo ThinkSystem servers are widely certified for running mission-critical SAP workloads, including SAP HANA and SAP S/4HANA. The fact that Lenovo's mainstream servers, such as the ThinkSystem SR630 V4, are certified for core SAP applications confirms their reliability, performance, and compatibility within the SAP landscape. This established relationship and certification portfolio make them a trusted hardware choice for running any critical SAP component, including SAP EIC.

3 Architectural overview

The core structure of SAP EIC is based on a Kubernetes-based container environment and comprises several key components.

The first is Edge Lifecycle Management, which takes care of onboarding, setup, upgrades, and monitoring of the software. The runtime and operations layer includes the controllers and worker nodes that deploy and execute integration flows and APIs, along with local security and artifact management. Supporting services such as PostgreSQL for persistence, Redis for caching, and a message broker like Solace provide the necessary backing infrastructure, while load balancers handle distribution of incoming traffic.

Connectivity and security are addressed through local authentication and authorization, certificate and key management, and adapters that link to systems via RFC, HTTP, SOAP, or Kafka. A secure connection to the SAP BTP cloud ensures that design and governance remain centralized while execution stays local. Finally, governance and monitoring components manage API proxies, logging, metrics, and policy enforcement, making sure integrations are both controlled and transparent.

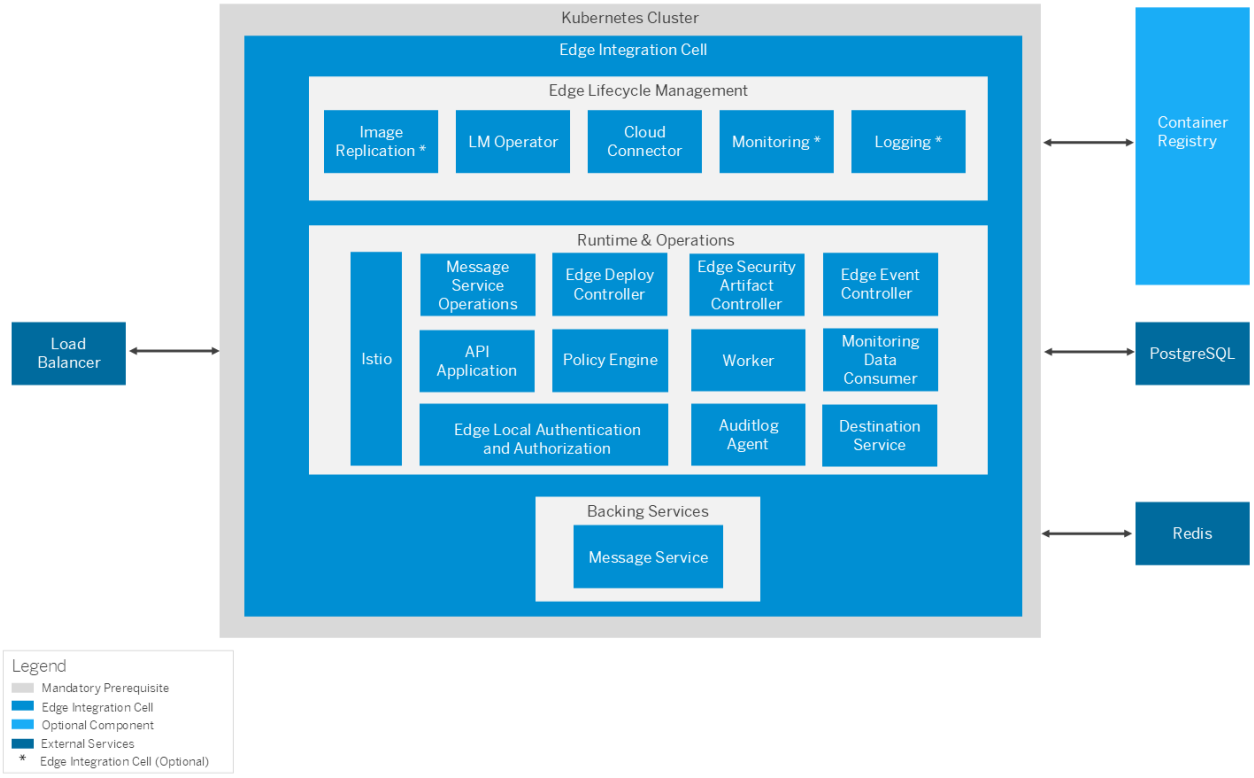

Figure 2: Edge Integration Cell and its components

Edge Integration Cell supports flexible deployment options to meet different computing scenarios:

- Baremetal – the container application platform runs directly on the operating system of the baremetal machine. This method ensures low latency processing for integration flows and full access to hardware resources with no virtualization overhead.
- Virtualized – the container application platform runs on an operating system that is inside a VM (for example VMware or KVM). This approach makes effective use of existing virtualization investments and physical resources like CPU and memory.

- Cloud based – runs on Kubernetes clusters provided by public cloud (Azure Kubernetes Service, Google Kubernetes Engine, Amazon Elastic Kubernetes Service). Elastic scalability, simplified management of Kubernetes clusters and cost optimization for workloads that do not need permanent edge deployment are the strong points of deploying Edge Integration Cell on cloud.

While Edge Integration Cell supports deployment across baremetal, virtualized, and cloud environments, the architecture described here focuses on the baremetal deployment.

The elements of a full stack baremetal architecture consisting of at least 3 Lenovo ThinkSystem servers is visualized in the figure below:

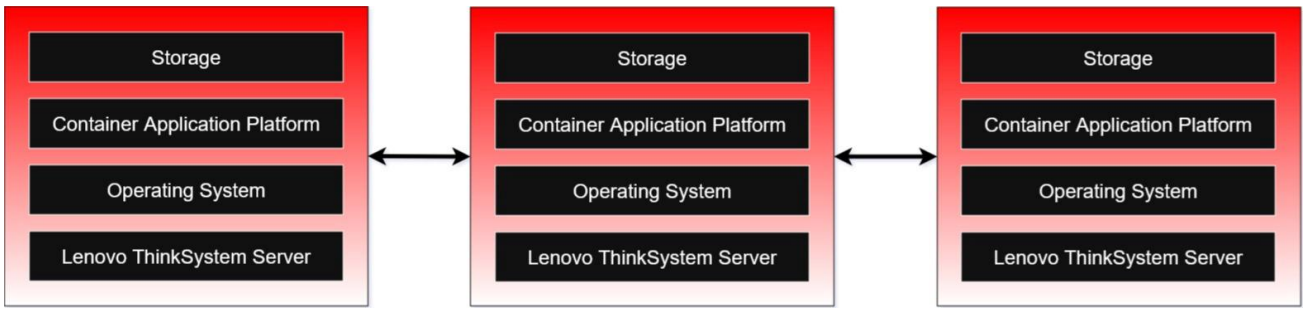

The choice of operating system is strongly related to the choice of container platform as some container platforms can only run on a small number of operating systems. For non-cloud-based deployments SAP Edge Integration Cell can run on SUSE Rancher Kubernetes Engine 2 (RKE2) and Red Hat OpenShift Container Platform (OCP). Red Hat OCP can run on Red Hat Enterprise Linux (RHEL) and Red Hat Enterprise Linux CoreOS (RHCOS). See https://access.redhat.com/articles/4763741 for more details.

SUSE RKE2 is supported on top of the following OS versions:

- SUSE Linux Enterprise Server (SLES)
- SUSE Linux Enterprise Micro (SLE micro)
- openSUSE Leap
- SUSE Liberty
- Oracle Linux
- Red Hat Enterprise Linux (RHEL)
- Rocky Linux
- Amazon Linux
- Ubuntu

See https://www.suse.com/suse-rke2/support-matrix/all-supported-versions/ for more details.

When deploying SAP Edge Integration Cell (EIC) on SUSE RKE2, persistent storage is required for stateful components like the messaging broker, integration runtime state, and databases. Solutions for persistent storage include:

- SUSE Longhorn – a lightweight, distributed block storage system designed for Kubernetes and edge environments. It is developed by SUSE and fully integrated with SUSE Rancher and RKE2.

It provides high availability even with a small number of nodes and works well on both baremetal and virtualized nodes. See https://longhorn.io/ for more details.

- NetApp Trident - an open-source Container Storage Interface (CSI) driver developed by NetApp for Kubernetes and containerized environments. It enables Kubernetes workloads to dynamically provision and manage persistent storage from NetApp storage systems. A good choice for on-prem SAP Edge Integration Cell deployments is ONTAP which is NetApp's flagship unified storage OS, available on AFF (All Flash FAS), FAS (Hybrid Flash), and ASA (All-Flash SAN Array) systems. See https://www.netapp.com/trident/ for more details.
- OpenShift Data Foundation - a software-defined storage solution for the Red Hat OpenShift Container Platform. It provides a unified data management and storage platform for persistent workloads and data services, designed to run natively on OpenShift. OpenShift Data Foundation is built on several open-source projects, primarily Ceph, Rook, and NooBaa. See https://www.redhat.com/en/resources/add-capabilities-enterprise-deployments-datasheet for more details.

The Lenovo solution discussed for the remainder of this document uses the following elements:

- Baremetal environment: Lenovo ThinkSystem SR630 V4
- Operating system: SUSE Linux Enterprise Micro
- Container application platform: SUSE Rancher Kubernetes Engine 2 (RKE2)
- Persistent storage: SUSE Longhorn

The reason for choosing these components is discussed in the following section.

4 Component model

4.1 Lenovo ThinkSystem V4

The Lenovo ThinkSystem V4 series provides a robust platform for hybrid cloud environments. Built on the Intel® Xeon® 6th Generation processors, these servers combine advanced processing power, flexible memory and storage options, and enterprise-class reliability. Within this series, the ThinkSystem SR630 V4 offer complementary capabilities for running mission-critical applications, databases, and analytics workloads in hybrid cloud architectures.

Lenovo ThinkSystem SR630 V4

The Lenovo ThinkSystem SR630 V4 is a 1U, two-socket rack server optimized for high-density deployments where performance, efficiency, and scalability are key. It is designed to deliver exceptional compute capability in a compact form factor, making it ideal for modern data centers, edge environments, and cloud workloads that demand strong performance per watt and per rack unit.

Key features include:

- Compact 1U design: Optimized for space-constrained environments and edge deployments where density and energy efficiency are critical.
- Memory scalability: Supports up to 32 DDR5 DIMMs, offering strong performance for mid-sized and virtualized workloads.
- Processor configuration: Supports dual Intel® Xeon® Scalable processors (P-cores) to deliver high single-threaded and multi-threaded performance in a smaller footprint.
- Streamlined storage: Supports up to 16x NVMe drives without oversubscription of PCIe lanes (1:1 connectivity) and without the need for additional NVMe adapters.

The SR630 V4 is ideal for mid-sized applications, distributed workloads, and scenarios where high-density, energy-efficient computing is a priority.

Figure 4: Lenovo SR630 V4

Together with other Lenovo ThinkSystem platforms, the SR630 V4 provides a flexible and cloud-ready infrastructure foundation for hybrid environments, enabling organizations to deploy critical workloads confidently across on-premises and cloud infrastructures with scalability, efficiency, and resilience.

4.2 SUSE Linux Enterprise Micro

SUSE Linux Enterprise Micro (SLE Micro) is a lightweight, immutable operating system purpose-built for running containerized and virtualized workloads across hybrid cloud and edge environments. As part of the SUSE Linux Enterprise family, SLE Micro delivers enterprise-grade security, lifecycle management, and support while maintaining a streamlined, minimal footprint optimized for cloud-native architectures.

At the foundation of SLE Micro is an immutable, read-only root file system. This ensures consistent deployments, eliminates configuration drift, and simplifies compliance. Updates are applied using transactional mechanisms: changes are atomic, verifiable, and can be rolled back if necessary. This approach provides predictable behavior and minimizes downtime—an essential capability for distributed hybrid and edge systems.

SLE Micro is specifically engineered to serve as a container host OS. It integrates seamlessly with container runtimes such as Podman and Docker and is certified for Kubernetes. Its lightweight design reduces resource overhead, making it highly efficient for running microservices-based architectures. By providing a secure and stable foundation, SLE Micro ensures that containers and microservices run consistently across environments.

SLE Micro integrates natively with SUSE Rancher, providing a robust and enterprise-ready platform for Kubernetes management. This integration enables:

- Centralized deployment and lifecycle management of Kubernetes clusters.
- Simplified orchestration of containers and microservices across hybrid and multi-cloud environments.
- Streamlined operations at scale, whether running in the core data center, public cloud, or at the edge.

Together, SLE Micro and Rancher form a powerful foundation for modern cloud-native environments, delivering consistency, reliability, and enterprise-grade support for Kubernetes-driven workloads.

Key Capabilities

- Container-optimized host: Purpose-built to run containerized applications efficiently and securely.
- Enterprise security: Includes SELinux, AppArmor, and built-in cryptography for strong workload isolation and data protection.
- Minimal footprint: Reduces attack surfaces and simplifies lifecycle management.
- Hybrid deployment flexibility: Suitable for bare-metal servers, virtual machines, and cloud instances, ensuring consistent operations across the hybrid landscape.
- Long-term support: Backed by SUSE with enterprise-grade updates, compliance, and technical support.

SUSE Linux Enterprise Micro (SLE Micro) delivers a secure, immutable, and container-optimized operating system ideally suited for running microservices in hybrid cloud and edge environments. Its native integration with Rancher and Kubernetes makes it a natural choice for enterprises adopting cloud-native architectures, ensuring scalability, simplified operations, and enterprise-grade reliability.

4.3 SUSE Rancher Kubernetes Engine 2 (RKE2)

SUSE RKE2 (Rancher Kubernetes Engine 2) is a next-generation, enterprise-grade Kubernetes distribution from SUSE. It is a fully conformant Kubernetes distribution that focuses on security and compliance. While it's often associated with the U.S. Federal Government sector (and was previously known as RKE Government), its robust design and features make it an ideal choice for any organization

that requires a highly secure and reliable Kubernetes platform, whether on-premises, in the cloud, or at the edge.

RKE2 is the successor to the original RKE (Rancher Kubernetes Engine), but it has a different architectural foundation. It combines the operational simplicity of SUSE's lightweight Kubernetes distribution, K3s, with a close alignment to upstream Kubernetes, providing a robust and production-ready platform for mission-critical workloads.

The core philosophy of RKE2 is to be "secure-by-default." This is achieved through several key architectural and feature choices:

- Security and Compliance: RKE2 is engineered to meet strict security requirements. By default, it is configured to pass the CIS Kubernetes Benchmark with minimal manual intervention. It also supports FIPS 140-2 compliance for cryptographic modules, which is a critical requirement for many government and highly regulated industries.
- Simplified Operations: RKE2 is designed to be easy to install and manage. Its single binary installation model simplifies deployment, and its configuration options are straightforward, using sensible defaults to reduce the burden on operators.
- No Docker Dependency: Unlike the original RKE, which relied on Docker, RKE2 uses containerd as its container runtime. This aligns with the direction of the upstream Kubernetes project and provides a more streamlined and efficient runtime environment.
- Static Pods for Control Plane: RKE2 runs the Kubernetes control plane components (like the API server, controller manager, and scheduler) as static pods managed by the kubelet process. This simplifies the management of the control plane and ensures its high availability.
- Wide Platform Support: RKE2 is not tied to a specific operating system. It can be deployed on a variety of Linux distributions, including SUSE Linux Enterprise Server (SLES) and others.
- Scalable Architecture: The design of RKE2, with its use of an etcd datastore and robust networking options, allows it to scale from small, single-node clusters to large, highly available deployments.

A typical RKE2 cluster consists of:

- Server Nodes (Control Plane): These nodes run the Kubernetes control plane components, including the kube-apiserver, kube-scheduler, and kube-controller-manager. They also run the etcd datastore, which stores the cluster's state. For high availability, it is recommended to have multiple server nodes.
- Agent Nodes (Worker Nodes): These nodes run the containerized applications. They run the kubelet agent, which manages the containers, and the kube-proxy, which handles networking rules.

Using SUSE RKE2 for SAP Edge Integration Cell is a well-established and validated approach. SAP officially supports and recommends specific Kubernetes distributions for Edge Integration Cell, and SUSE RKE2 was the first and is the most commonly used validated option.

4.4 SUSE Longhorn

SUSE Longhorn is an open-source, cloud-native distributed block storage solution for Kubernetes. Designed to be lightweight, reliable, and easy to use, it provides persistent storage for stateful applications running in Kubernetes clusters. Unlike traditional storage solutions that require dedicated, external hardware, Longhorn is built to run directly within your Kubernetes cluster, leveraging the existing local storage on your nodes.

Longhorn's strength lies in its comprehensive set of enterprise-grade storage features that are delivered in a simple, cloud-native way.

- Distributed and Highly Available: Longhorn provides high availability by synchronously replicating data across multiple nodes in the cluster. If a node or disk fails, your data remains accessible, and Longhorn automatically rebuilds the lost replicas to restore redundancy.
- Snapshots and Backups: It offers built-in features for creating point-in-time snapshots and backups. Snapshots are lightweight and instant, while backups can be stored on external, S3-compatible object storage (like AWS S3 or MinIO) or on an NFS share, providing a robust disaster recovery plan.
- Incremental Backups: Longhorn performs incremental backups, meaning it only copies the data that has changed since the last backup. This dramatically reduces the amount of storage space and network bandwidth required for data protection.
- Intuitive UI and API: Longhorn includes a clean, web-based dashboard that simplifies the management of volumes, nodes, and backups. It also provides a comprehensive API, allowing for complete automation of storage tasks.
- Volume Live Migration: You can non-disruptively move a volume from one node to another, a critical feature for planned maintenance or load rebalancing without affecting running applications.
- Dynamic Provisioning: Longhorn is a Container Storage Interface (CSI) compliant solution. This allows it to dynamically provision Persistent Volumes (PVs) when an application requests storage via a Persistent Volume Claim (PVC), automating the entire storage provisioning workflow.

Longhorn's architecture is built on the principle of a "microservice per volume.":

- Longhorn Manager: A Kubernetes DaemonSet that runs on every node, acting as the control plane. It is responsible for creating and managing volumes, handling API calls, and orchestrating the entire system.
- Longhorn Engine: A dedicated storage controller that runs for each individual volume. This is a lightweight, single-process container that handles all I/O operations for its specific volume. This design ensures that a problem with one volume does not affect the performance or availability of others.
- Longhorn Replicas: These are the actual data components of a volume. Each replica is a separate process that stores a copy of the volume's data. Replicas are strategically placed on different nodes to provide data redundancy.

This microservice-based architecture is what gives Longhorn its lightweight, resilient, and non-disruptive nature.

Longhorn is a versatile solution that is particularly well-suited for:

- Edge and IoT Deployments: Its lightweight design and ability to run on commodity hardware make it perfect for edge locations with limited resources and intermittent connectivity.
- On-Premises Kubernetes Clusters: It provides a cost-effective alternative to expensive, proprietary external storage arrays, leveraging your existing hardware investments.
- DevOps and Development Teams: Its easy-to-use interface, dynamic provisioning, and robust API make it simple for developers to manage their own storage, accelerating application development cycles.
- Anyone Requiring Enterprise-Grade Data Protection: With its snapshots, backups, and disaster recovery features, Longhorn is a strong choice for mission-critical, stateful applications like databases (e.g., PostgreSQL, MySQL), message queues, and other persistent workloads.

In essence, SUSE Longhorn democratizes enterprise-grade persistent storage, making it accessible and manageable for a wide range of Kubernetes users without the complexity or cost of traditional storage solutions.

The combination of SUSE Longhorn and SAP Edge Integration Cell provides a secure, highly available, and operationally efficient solution for running mission-critical SAP workloads at the edge. SUSE provides detailed best practice guides and documentation that outline how to install and configure Longhorn in enterprise environments. This eliminates guesswork and reduces the risk of deployment, ensuring you have a production-ready and supported setup.

Using SUSE Linux Enterprise (SLE) Micro, SUSE RKE2, and SUSE Longhorn together for SAP Edge Integration Cell (EIC) provides an integrated, edge-optimized, secure, and fully open-source stack that has been validated by both SUSE and SAP.

4.5 Edge Integration Cell worker node

From an operational perspective, EIC is deployed as a containerized runtime, typically on Kubernetes or comparable orchestration platforms (RKE2, Red Hat OpenShift, etc.). It is provisioned, configured, and managed centrally through SAP Integration Suite, ensuring consistency of lifecycle management, governance, and monitoring. Once deployed, it executes integration flows locally, enabling secure, near real-time data exchange across cloud applications, on-premise systems, and edge environments.

The interaction between SAP Integration Suite and EIC is visualized in

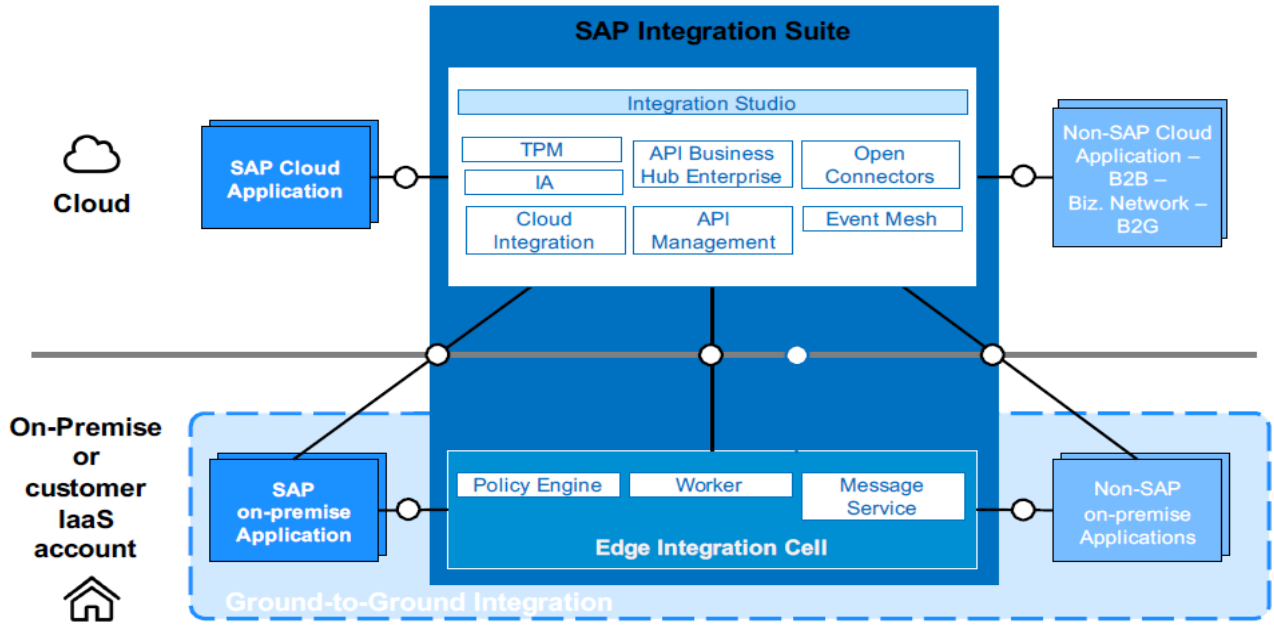

Figure 5 below.

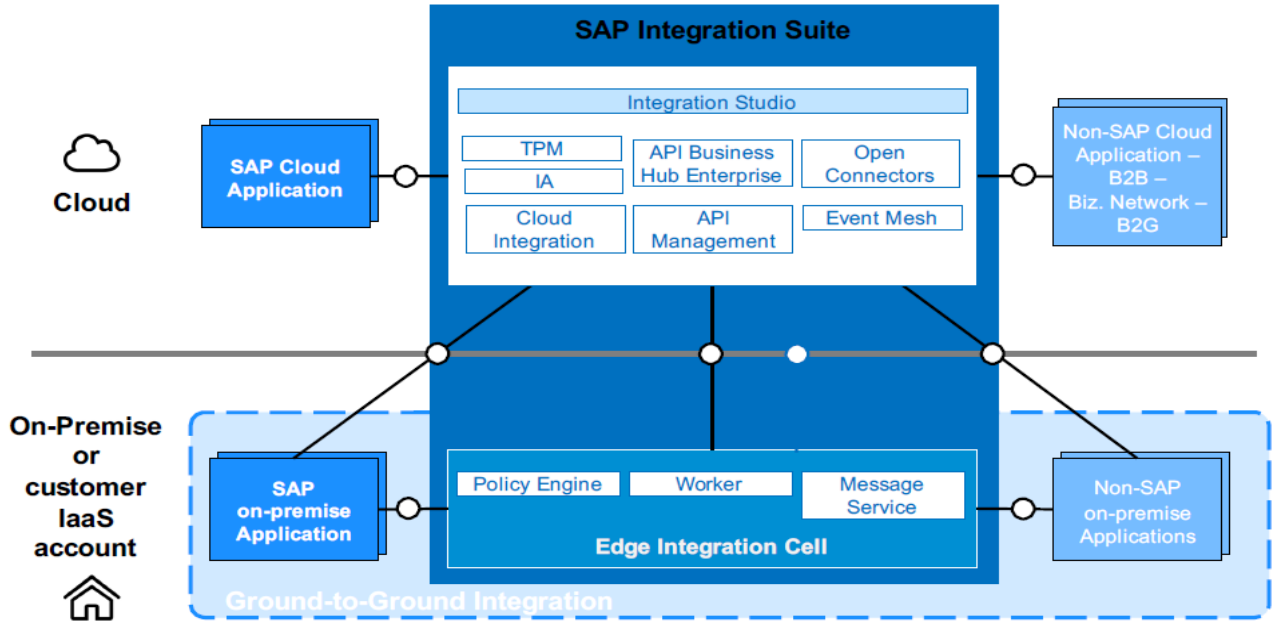

Figure 5: Interaction of SAP Integration Suite and EIC

4.6 External database (PostgreSQL)

PostgreSQL, often called Postgres is a powerful, open-source relational database management system (RDBMS) known for its robustness, standards compliance, and extensibility. It supports advanced SQL features, complex queries, and transactional consistency, making it a reliable platform for enterprise-grade applications. PostgreSQL is designed to manage structured, long-lived data with a strong focus on durability, data integrity, and security.

PostgreSQL is widely used across industries for:

- Enterprise application backends requiring reliable relational data storage.
- Transactional systems where ACID (Atomicity, Consistency, Isolation, Durability) compliance is critical.
- Analytics and reporting on structured datasets.
- Integration with modern workloads including microservices, geospatial applications, and AI/ML.
- Multi-model data handling (support for JSON, XML, key-value storage alongside relational tables).

Its scalability and extensibility make PostgreSQL a trusted backbone for mission-critical applications.

In the SAP Edge Integration Cell (EIC) architecture, PostgreSQL serves as the persistent data store that complements the in-memory caching capabilities of Redis:

- Structured data storage: PostgreSQL securely stores structured and transactional data generated by integration flows.
- Durability and compliance: It provides reliable and secure storage of business-critical information, helping meet regulatory and data residency requirements in hybrid environments.
- Consistency for integration pipelines: By maintaining ACID properties, PostgreSQL guarantees data integrity for the integration processes executed within EIC.
- Persistent and reliable storage: PostgreSQL ensures that important state and metadata are stored permanently and remain accessible for governance, compliance, and audit purposes.

In summary, PostgreSQL provides the enterprise-grade persistence layer for SAP EIC, ensuring structured data is managed reliably, securely, and in accordance with business and regulatory requirements. Its durability and transactional consistency make it a cornerstone of the runtime, supporting the performance, reliability, and compliance needs of hybrid integration deployments.

4.7 External datastore (Redis)

Redis (Remote Dictionary Server) is an open-source, in-memory data structure store widely used as a cache, message broker, and key–value database. It is designed for high performance and low latency, capable of processing millions of operations per second. Unlike traditional relational databases, Redis stores data directly in memory, making it extremely fast for lookups, updates, and temporary data storage.

Redis is commonly used across modern IT architectures to:

- Cache frequently accessed data to reduce load on backend systems.
- Store transient data such as session information, temporary states, and application metadata.
- Provide fast key–value storage for microservices and distributed systems.
- Enable high-throughput messaging and event-driven architectures.

By handling data that does not require long-term persistence, Redis accelerates application performance and reduces latency in end-to-end transactions.

In the SAP Edge Integration Cell (EIC) architecture, Redis plays a complementary role alongside the PostgreSQL database:

- Redis for caching: Redis is used to store frequently accessed or temporary data in-memory, improving the performance of integration flows by reducing the need for repeated queries to the persistent database.
- PostgreSQL for persistence: PostgreSQL acts as the relational database for structured, long-lived data that must be stored and managed reliably.
- Optimized pipeline: By offloading transient workloads to Redis and reserving PostgreSQL for durable storage, SAP EIC achieves both speed and reliability in its runtime pipeline.
- Support for resilience: In hybrid landscapes with high transaction volumes, Redis helps minimize latency, ensuring that integration flows execute efficiently even under heavy load.

Redis is a critical component of the SAP EIC unified runtime, delivering high-speed caching and transient data handling that complements PostgreSQL's role as the persistent data store. Together, Redis and PostgreSQL ensure that SAP EIC can execute integration flows with both the performance required at the edge and the data integrity needed for compliance and governance.

4.8 Monitoring (Prometheus)

Prometheus is an open-source monitoring and alerting toolkit that has become a foundational component in the cloud-native landscape, particularly with Kubernetes. Its primary function is to collect and store real-time metrics as time-series data using a flexible, multi-dimensional data model. The core mechanism is a pull model, where the central Prometheus server periodically "scrapes" metrics from instrumented application and service endpoints over HTTP. It features a powerful query language called PromQL for selecting and aggregating this data for analysis, and it uses a separate Alertmanager component to handle and route notifications based on predefined alerting rules. While Prometheus itself offers basic graphing

While Prometheus has a basic web interface for querying and graphing, it is most used in conjunction with Grafana, an open-source data visualization and analytics platform. Grafana can use Prometheus as a data source to create rich, interactive dashboards that provide a real-time view of system health and performance.

Prometheus in SAP EIC is used as the monitoring and metrics collection backbone. Since EIC runs inside a customer-managed Kubernetes environment, it needs a reliable way to observe the health, performance, and resource usage of all its microservices, integration runtimes, and supporting components.

Prometheus scrapes metrics from the various EIC components—such as the runtime pods, worker nodes, messaging services, and supporting databases like PostgreSQL or Redis—and stores them in a time-series database. These metrics cover aspects like CPU and memory consumption, request counts, latency, error rates, and availability.

The collected data is then used in multiple ways: operators can query Prometheus directly, integration with Grafana provides dashboards, and the metrics can trigger alerts when thresholds are crossed. This makes Prometheus essential for day-to-day operations of SAP EIC because it ensures administrators can detect issues early, troubleshoot bottlenecks, and maintain service-level health of the integration cell.

5 Deployment considerations

We recommend using three to six worker nodes depending on workload requirements. A reasonable test environment usually comprises three nodes, while a production cluster starts with three nodes and should follow SAP and SUSE EIC sizing guidelines (SAP Note 3705131).

As a starting point we have provided three bill of materials for small, medium and large environments. They should be considered a reasonable minimum, it is of course best practices to increase a component according to specific use case needs.

If you wish to follow a certain company policy and leverage components from a specific vendor you can replace these – for example network cards – with other supported options from the Lenovo portfolio. There

5.1 BOM for small EIC worker node

Part number	Product description	Qty
7DG9CTO1WW	ThinkSystem SR630 V4-3yr Base Warranty	1
C1XE	ThinkSystem 1U V4 10x2.5" Chassis	1
C3J9	ThinkSystem General Computing - Max Performance	1
BVGL	Data Center Environment 30 Degree Celsius / 86 Degree Fahrenheit	1
C5RD	Intel Xeon 6515P 16C 150W 2.3GHz Processor	2
C1XJ	ThinkSystem 1U V4 Performance Heatsink	2
C0U9	ThinkSystem 32GB TruDDR5 6400MHz (1Rx4) RDIMM	8
C1WM	ThinkSystem 2.5" U.2 PM9D5a 800GB Mixed Use NVMe NVMe PCIe 5.0 x4 HS SSD	5
C21X	ThinkSystem 1U V4 10x2.5" NVMe Gen5 Backplane	1
C0JJ	ThinkSystem M.2 RAID B540p-2HS SATA/NVMe Adapter	1
CBSZ	ThinkSystem M.2 VA 480GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	2
BCD4	ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port OCP Ethernet Adapter	2
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
C1YT	ThinkSystem 1U V4 Performance Fan Module	4
C2DH	ThinkSystem Toolless Slide Rail Kit V4	1
BPKR	TPM 2.0	1
B7XZ	Disable IPMI-over-LAN	1
BE0F	N+N Redundancy Without Over-Subscription	1
SCY0	Lenovo XClarity XCC3 premier - FOD	1
1340	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1

5.2 BOM for medium EIC worker node

Part number	Product Description	Qty
7DG9CTO1WW	ThinkSystem SR630 V4-3yr Base Warranty	1
C1XE	ThinkSystem 1U V4 10x2.5" Chassis	1
C3J9	ThinkSystem General Computing - Max Performance	1
BVGL	Data Center Environment 30 Degree Celsius / 86 Degree Fahrenheit	1
C5QT	Intel Xeon 6530P 32C 225W 2.3GHz Processor	2
C1XJ	ThinkSystem 1U V4 Performance Heatsink	2
BYTJ	ThinkSystem 32GB TruDDR5 6400MHz (2Rx8) RDIMM	16
C0ZR	ThinkSystem 2.5" U.2 VA 1.6TB Mixed Use NVMe PCIe 5.0 x4 HS SSD	5
C21X	ThinkSystem 1U V4 10x2.5" NVMe Gen5 Backplane	1
C0JJ	ThinkSystem M.2 RAID B540p-2HS SATA/NVMe Adapter	1
CBSZ	ThinkSystem M.2 VA 480GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	2
BCD4	ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port OCP Ethernet Adapter	2
C1YH	ThinkSystem SR630 V4 x16/x16 PCIe Gen5 Cable Riser 1	1
C1Z7	ThinkSystem SR630 V4 Full Height+Low Profile Riser1 Cage	1
C0U3	ThinkSystem 2000W 230V Titanium CRPS Premium Hot-Swap Power Supply	2
C1YT	ThinkSystem 1U V4 Performance Fan Module	4
C2DH	ThinkSystem Toolless Slide Rail Kit V4	1
BPKR	TPM 2.0	1
B7XZ	Disable IPMI-over-LAN	1
BE0E	N+N Redundancy With Over-Subscription	1
SCY0	Lenovo XClarity XCC3 premier - FOD	1
1340	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1

5.3 BOM for large EIC worker node

Part number	Product Description	Qty
7DG9CTO1WW	ThinkSystem SR630 V4-3yr Base Warranty	1
C1XE	ThinkSystem 1U V4 10x2.5" Chassis	1
C3J9	ThinkSystem General Computing - Max Performance	1
BVGL	Data Center Environment 30 Degree Celsius / 86 Degree Fahrenheit	1
C5R3	Intel Xeon 6740P 48C 270W 2.1GHz Processor	2
C1XJ	ThinkSystem 1U V4 Performance Heatsink	2
BYTJ	ThinkSystem 32GB TruDDR5 6400MHz (2Rx8) RDIMM	24
C0ZQ	ThinkSystem 2.5" U.2 VA 3.2TB Mixed Use NVMe PCIe 5.0 x4 HS SSD	4
C21X	ThinkSystem 1U V4 10x2.5" NVMe Gen5 Backplane	1

C0JJ	ThinkSystem M.2 RAID B540p-2HS SATA/NVMe Adapter	1
CBSZ	ThinkSystem M.2 VA 480GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	2
BCD4	ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port OCP Ethernet Adapter	2
C1YH	ThinkSystem SR630 V4 x16/x16 PCIe Gen5 Cable Riser 1	1
C1Z7	ThinkSystem SR630 V4 Full Height+Low Profile Riser1 Cage	1
C0U3	ThinkSystem 2000W 230V Titanium CRPS Premium Hot-Swap Power Supply	2
C1YT	ThinkSystem 1U V4 Performance Fan Module	4
C2DH	ThinkSystem Toolless Slide Rail Kit V4	1
BPKR	TPM 2.0	1
B7XZ	Disable IPMI-over-LAN	1
BE0E	N+N Redundancy With Over-Subscription	1
SCY0	Lenovo XClarity XCC3 premier - FOD	1
1340	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1

5.4 Network configuration

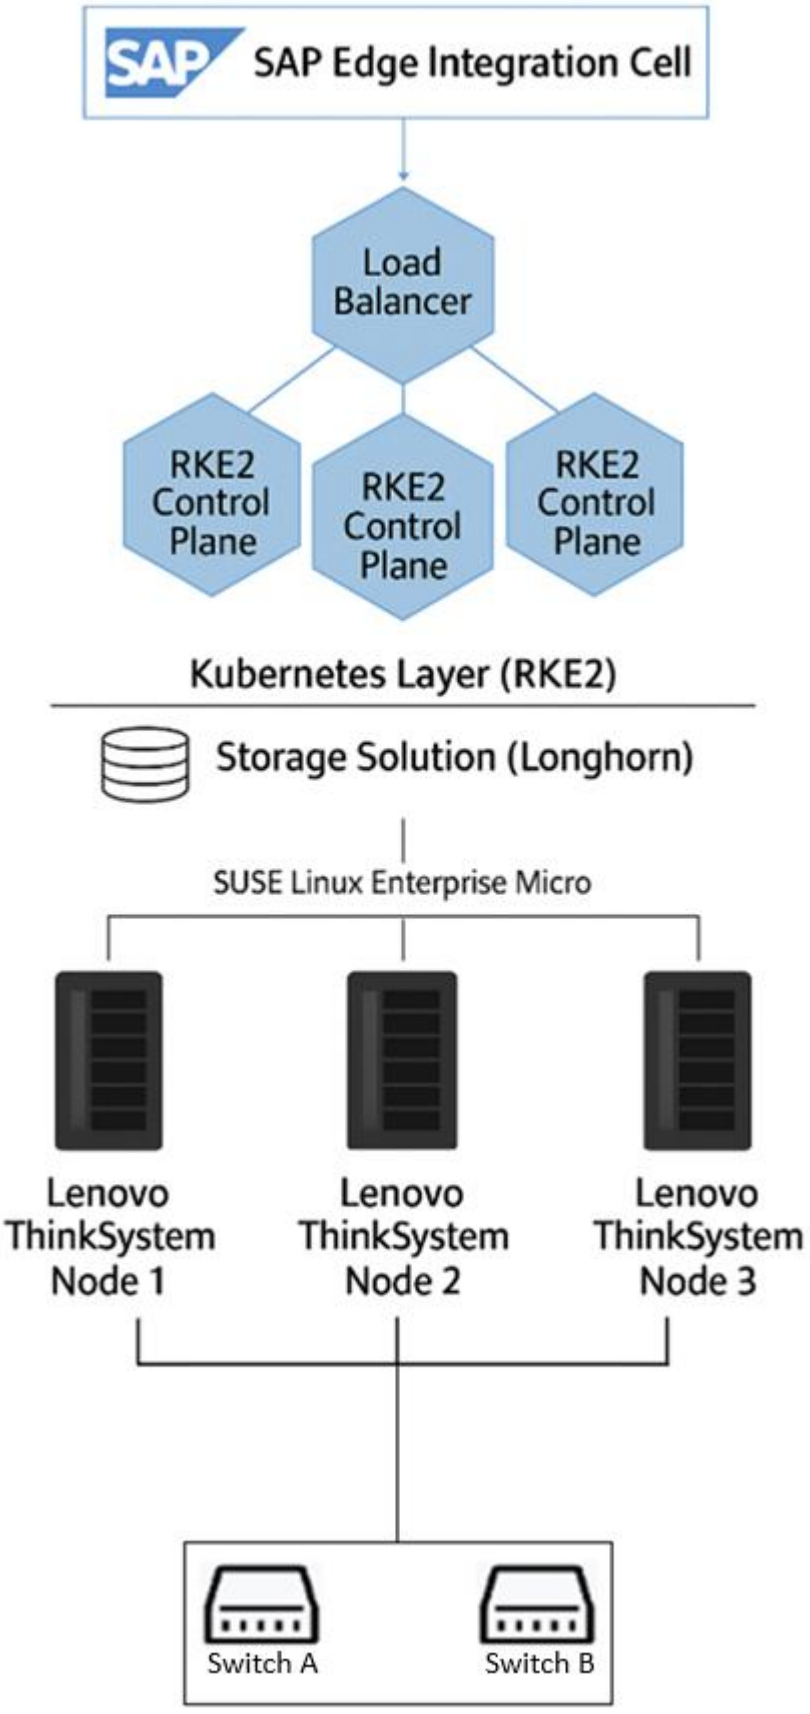

Figure 6: EIC mapped onto Lenovo architecture

The diagram presents the recommended network deployment for SAP Edge Integration Cell (EIC) in an enterprise edge environment. It illustrates a high-availability RKE2 Kubernetes cluster deployed across three Lenovo ThinkSystem nodes, each running SUSE Linux Enterprise Micro as the host operating system. Together, these nodes form a robust and secure platform for the EIC workload, supported by SUSE Longhorn for distributed and resilient storage across the cluster.

Network connectivity is designed for redundancy and fault tolerance, featuring a bonded link (LACP) from each Lenovo ThinkSystem node to two independent switches. This configuration ensures continuous network availability and minimizes the risk of service interruption due to hardware or link failure.

A load balancer provides traffic distribution and high availability for the RKE2 control plane, while the upper layers represent the SAP Edge Integration Cell and its integration with the broader SAP ecosystem. This recommended topology enables secure, resilient, and production-ready deployment of SAP EIC in edge or on-premises environments, aligning with enterprise best practices for performance, scalability, and operational continuity.

5.5 Best practices

5.5.1. System and Infrastructure Design

Node Sizing and Hardware Selection: Use enterprise-grade Lenovo ThinkSystem servers with redundant power supplies and network interfaces. Size CPU, memory, and storage resources based on SAP EIC workload requirements and SAP’s sizing recommendations.

High Availability (HA): Deploy a minimum of three RKE2 control plane nodes to achieve HA. Configure SUSE Longhorn with at least three replicas for resilient, distributed storage across the nodes.

Network Redundancy: Connect each node via a bonded LACP link to two independent switches for fault-tolerant network connectivity and throughput optimization.

5.5.2. Operating System and Cluster Configuration

Operating System: Run SUSE Linux Enterprise Micro on all cluster nodes for a minimal, secure, and immutable OS foundation. Regularly apply SUSE-provided updates and security patches.

Cluster Configuration: Align the RKE2 version with SUSE Rancher’s supported versions. Enable certificate rotation, set up etcd snapshots (through RKE2), and verify automatic recovery processes.

Time Synchronization: Synchronize all nodes using NTP to maintain consistent timestamps for cluster components, logs, and certificate operations.

5.5.3. Storage and Data Management

SUSE Longhorn: Configure Longhorn as the primary persistent storage backend for SAP EIC workloads. Enable volume replication ($\geqslant$3) to ensure data availability in case of node failures.

Backups and Recovery: Schedule automated backups of persistent volumes and periodically test restore operations to verify data integrity.

5.5.4. Database Configuration (PostgreSQL and Redis)

PostgreSQL Database:

Deploy PostgreSQL as a highly available instance, either within the RKE2 cluster (as a StatefulSet) or externally as a managed service.

- Use persistent volumes provisioned by Longhorn for database storage.
- Enable regular backups, point-in-time recovery, and SSL/TLS for secure communication.
- Monitor database performance (I/O latency, connection pool usage) using integrated observability tools.

Redis Cache:

- Deploy it in replication mode or use Redis Sentinel for automatic failover.
- Ensure Redis data directories are backed by persistent volumes to preserve state in HA scenarios.
- Apply authentication and TLS to secure inter-component communication.
- Regularly monitor Redis metrics (memory usage, key eviction, replication status).

5.5.5. Security and Compliance

Authentication and Access Control: Integrate with enterprise identity providers (LDAP, Active Directory, or SAML) for centralized authentication.

Network Policies: Apply Kubernetes NetworkPolicies to restrict communication between namespaces and isolate Redis, PostgreSQL, and EIC workloads.

Certificates: Use cert-manager for automated TLS certificate provisioning and renewal to secure ingress and service-to-service traffic.

5.5.6. Monitoring and Maintenance

Observability Stack: Enable cluster monitoring through Prometheus and Grafana or the Rancher integrated monitoring stack. Include dashboards for Redis, PostgreSQL, and Longhorn.

Logging: Aggregate logs centrally using Fluent Bit, Loki, or a compatible log management system to support auditing and troubleshooting.

Lifecycle Management: Perform rolling updates using SUSE Rancher to ensure minimal downtime and maintain consistency across nodes and workloads.

5.5.7. Performance and Scalability

Resource Tuning: Define CPU and memory requests/limits for EIC workloads, Redis, and PostgreSQL to ensure predictable performance and prevent contention.

Scaling Strategy: Add worker nodes horizontally as workload demands increase. Validate storage replication and network performance before scaling.

Load Balancing: Use a Layer 4 or Layer 7 load balancer (hardware or cloud-based) to manage EIC control-plane and ingress traffic efficiently.

6 Summary

SAP Edge Integration Cell (EIC) provides a containerized runtime for executing integration flows locally, ensuring low-latency performance and compliance with data residency requirements. For sensitive or mission-critical processes that must remain on-premise, Lenovo ThinkSystem servers offer a robust, high-performance platform to run EIC reliably. This combination allows organizations to harness the agility of the cloud while keeping critical workloads secure within their own datacenter.

For further information, please visit:

Edge Integration Cell

- Hybrid Deployment Using Edge Integration Cell
- SAP Integration Suite

SUSE

- SUSE Linux Micro
- SUSE Rancher Prime: RKE2
- SUSE Storage (Longhorn)
- SAP Edge Integration Cell on SUSE

Others

- Prometheus
- PostgreSQL
- Redis

Document History

Version 1.0	January 2026	Initial version