

Layered Stack Deployment of Rancher Kubernetes Engine Government

Integrated with Dell Technologies (R)

Layered Stack Deployment of Rancher Kubernetes Engine Government: Integrated with Dell Technologies (R)

SUSE Linux Enterprise Server 15 SP3, Rancher Kubernetes Engine Government 1.20.14

The purpose of this document is to provide an overview and procedure of implementing SUSE (R) and partner offerings for Rancher Kubernetes Engine Government (RKE2), a Kubernetes distribution that runs entirely within containers on bare-metal and virtualized nodes. RKE2 solves the problem of installation complexity and the operation is both simplified and easily automated, while entirely accommodating the operating system and platform it is running on. Also being a hardened, FIPS-enabled version, it adopts a compliance-based approach toward security, targeting standard risk management frameworks and best practices with the goal of stronger defense for cloud-native applications.

Disclaimer: Documents published as part of the series SUSE Technical Reference Documentation have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

Publication Date: 2022-03-28

<https://documentation.suse.com> 

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Scope	2
1.3	Audience	2
2	Business aspect	3
2.1	Business problem	3
2.2	Business value	4
3	Architectural overview	7
3.1	Solution architecture	7
4	Component model	9
4.1	Component overview	9
4.2	Software - Rancher Kubernetes Engine Government	9
4.3	Software - SUSE Linux Enterprise Server	12
4.4	Compute Platform	12
	Dell EMC PowerEdge R640 Rack Servers	13 • Dell EMC PowerEdge R740 Rack Servers
	Dell EMC PowerEdge R650 Rack Servers	18 • Dell EMC PowerEdge R750 Rack Servers
		20
5	Deployment	23
5.1	Deployment overview	23
5.2	Compute Platform	24
5.3	SUSE Linux Enterprise Server	24
5.4	Rancher Kubernetes Engine Government	26

6	Summary	32
7	References	33
8	Glossary	35
9	Appendix	38
9.1	Compute platform bill of materials	38
9.2	Software bill of materials	43
9.3	Documentation configuration / attributes	45
10	Legal Notice	46
11	GNU Free Documentation License	47

1 Introduction

On the digital transformation journey to a full cloud-native landscape, the use of microservices becomes the main approach with the dominant technology for such container orchestration being Kubernetes.¹ With its large community of developers and abundant features and capabilities, Kubernetes has become the de-facto standard and is included across most container-as-a-service platforms. With all of these technologies in place, both developer and operation teams can effectively deploy, manage and deliver functionality to their end users in a resilient and agile manner.

1.1 Motivation

Once on such a digital transformation journey, also relevant to focus on areas like:

Workload(s)

Determine how to manage and launch internally developed containerized, microservice workloads

Kubernetes

As developers and organizations continue their journey from simple, containerized microservices toward having these workloads orchestrated and deployed where ever they need, being able to install, monitor and use such Kubernetes infrastructures is a core need. Such deployments, being Cloud Native Computing Foundation (CNCF²) conformant and certified³ are essential for both development and production workloads.

- With core focus on security and compliance, Rancher Kubernetes Engine Government inherits close alignment with upstream Kubernetes and provide usability, ease-of-operations, and deployment model for core use cases.


Compute Platform(s)

To optimize availability, performance, scalability and integrity, assess current system or hosting platforms

¹ <https://kubernetes.io/> ↗

² <https://www.cncf.io/> ↗

³ <https://www.cncf.io/certification/software-conformance> ↗

from Independent Hardware Vendors (IHV), such as [Dell Technologies \(https://www.dell.com/en-us\)](https://www.dell.com/en-us)  ® as the platform for physical, bare metal, hypervisors and virtual machines

1.2 Scope

The scope of this document is to provide a layered *reference configuration* for Rancher Kubernetes Engine Government. This can be done in a variety of scenarios to create an enterprise Kubernetes cluster deployment anywhere to provide a very secure environment.

1.3 Audience

This document is intended for IT decision makers, architects, system administrators and technicians who are implementing a flexible, software-defined Kubernetes platform. One should still be familiar with the traditional IT infrastructure pillars — networking, computing and storage — along with the local use cases for sizing, scaling and limitations within each pillars' environments.

2 Business aspect

Agility is driving developers toward more cloud-native methodologies that focus on microservices architectures and streamlined workflows. Container technologies, like Kubernetes, embody this agile approach and help enable cloud-native transformation.

By unifying IT operations with Kubernetes, organizations realize key benefits like increased reliability, improved security and greater efficiencies with standardized automation. Therefore, Kubernetes infrastructure platforms are adopted by enterprises to deliver:

Cluster Operations

Improved Production and DevOps efficiencies with simplified cluster usage and robust operations

Security Policy & User Management

Consistent security policy enforcement plus advanced user management on any Kubernetes infrastructure

Access to Shared Tools & Services

A high level of reliability with easy, consistent access to a broad set of tools and services

2.1 Business problem

Many organizations are deploying Kubernetes clusters everywhere — in the cloud, on-premises, and at the edge — to unify IT operations. Such organizations can realize dramatic benefits, including:

- Consistently deliver a high level of reliability on any infrastructure
- Improve DevOps efficiency with standardized automation
- Ensure enforcement of security policies on any infrastructure

However, simply relying on upstream Kubernetes alone can introduce extra overhead and risk because Kubernetes clusters are typically deployed:

- Without central visibility
- Without consistent security policies
- And must be managed independently

Deploying a scalable kubernetes infrastructure requires consideration of a larger ecosystem, encompassing many software and infrastructure components and providers. Further, the ability to continually address the needs and concerns of:

Developers

For those who focus on writing code to build their apps securely using a preferred workflow, providing a simple, push-button deployment mechanism of their containerized workloads where needed.

IT Operators

General infrastructure requirements still rely upon traditional IT pillars are for the stacked, underlying infrastructure. Ease of deployment, availability, scalability, resiliency, performance, security and integrity are still core concerns to be addressed for administrative control and observability.

Beyond the core infrastructure software layers of managed Kubernetes clusters, organizations may be also be impacted by:

Compute Platform

Potential inconsistencies and impacts of multiple target system platforms for the distributed deployments of the cluster elements, across:

- physical, baremetal, hypervisors and virtual machines

2.2 Business value

With Rancher Kubernetes Engine Government, the operation of Kubernetes is easily automated and entirely independent of the operating system and platform running. Using a supported version of the container runtime engine, one can deploy and run Kubernetes with Rancher Kubernetes Engine Government. It builds a cluster from a single command in a few minutes, and its declarative configuration makes Kubernetes upgrades atomic and safe.

By allowing operation teams to focus on infrastructure and developers to deploy code the way they want too, SUSE and the Rancher offerings helps bring products to market faster and accelerate an organization's digital transformation.

SUSE Rancher is a fundamental part of the complete software stack for teams adopting containers. It provides DevOps teams with integrated tools for running containerized workloads while also addressing the operational and security challenges of managing multiple Kubernetes clusters across any targeted infrastructure.

Developers


SUSE Rancher makes it easy to securely deploy containerized applications no matter where the Kubernetes infrastructure runs -- in the cloud, on-premises, or at the edge. Using Helm or the App Catalog to deploy and manage applications across any or all these environments, ensuring multi-cluster consistency with a single deployment process.

IT Operators

SUSE Rancher not only deploys and manages production-grade Kubernetes clusters from datacenter to cloud to the edge, it also unites them with centralized authentication, access control and observability. Further, it streamlines cluster deployment on bare metal or virtual machines and maintains them using defined security policies.

With this increased consistency of the managed Kubernetes infrastructure clusters, organizations benefit from an even higher level of the Cloud Native Computing model where each layer only relies upon the API and version of the adjacent layer, such as:

Compute Platform

Using the above software application and technology solutions with the server platforms offered by [Dell Technologies \(https://www.dell.com/en-us\)](https://www.dell.com/en-us)  provides many alternative for scale, cost-effectiveness and performance options that could align with local IT staff platform preferences:

- density-optimized - high performance and efficiency for big data and the most demanding workloads
- mission-critical - systems of intelligence to fuel your digital transformation in a world where time and data are the new currency and business continuity is expected
- composable - fully adaptable and ready for Hybrid-IT to future-proof your data center for today's workloads and tomorrow's disruptors
- IoT - realize the potential of the Internet of Things to provide compute at the network edge

- cloud - high-capacity, mass-compute open infrastructure with security and software to match
- and virtualized use cases.

3 Architectural overview

This section outlines the core elements of the Rancher Kubernetes Engine Government solution, along with the suggested target platforms and components.

3.1 Solution architecture

The figure below illustrates the high-level architecture overview of Kubernetes components on instances like Rancher Kubernetes Engine Government:




FIGURE 3.1: ARCHITECTURE OVERVIEW - RANCHER KUBERNETES ENGINE GOVERNMENT

A Kubernetes cluster consists of a set of nodes machines, called workers or agents, that host and run containerized applications in Pods. Every cluster has at least one worker node. The control plane manages the worker nodes and the Pods in the cluster. The provider API is a generic element that allows external interaction with the Kubernetes cluster.

Control Plane Components

The control plane's components make global decisions about the cluster (for example, scheduling), as well as detecting and responding to cluster events.

- kube-apiserver
 - The API server is a component of the Kubernetes control plane that exposes the Kubernetes API
- etcd

- Consistent and highly-available key value store used as Kubernetes' backing store for all cluster data.
- kube-scheduler
 - Control plane component that watches for newly created Pods with no assigned node, and selects a node for them to run on.
- kube-controller-manager
 - Control plane component that runs controller processes.

Node Components

Node components run on every node, maintaining running pods and providing the Kubernetes runtime environment.

- kubelet
 - An agent that runs on each node in the cluster. It makes sure that containers are running in a Pod.
- kube-proxy
 - A network proxy that runs on each node in your cluster, implementing part of the Kubernetes Service concept.

While all Rancher Kubernetes Engine Government roles can be installed on a single system, for the best availability, performance and security, the recommended deployment of a Rancher Kubernetes Engine Government cluster is a pair of nodes for the control plane role, at least three etcd role-based nodes and three or more worker nodes.



Note

Regardless of the deployment instance, Rancher Kubernetes Engine Government could always be deployed by SUSE Rancher or imported as a managed, downstream cluster.

4 Component model

This section describes the various components being used to create a Rancher Kubernetes Engine Government solution deployment, in the perspective of top to bottom ordering. When completed, the Rancher Kubernetes Engine Government instance can be used as the application infrastructure for cloud-native workloads and can be imported into SUSE Rancher for management.

4.1 Component overview

By using:

- Kubernetes Platform - Rancher Kubernetes Engine Government
- Operating System - SUSE Linux Enterprise Server
- Compute Platform
 - Dell EMC PowerEdge

you can create the necessary infrastructure and services. Further details for these components are described in the following sections.

4.2 Software - Rancher Kubernetes Engine Government

Rancher Kubernetes Engine Government also known as RKE2, is Rancher's next-generation Kubernetes distribution. It is a fully conformant Kubernetes distribution that focuses on security and compliance within the U.S. Federal Government sector. It solves the common frustration of installation complexity with Kubernetes by removing most host dependencies and presenting a stable path for deployment, upgrades, and rollbacks.

To meet these goals, Rancher Kubernetes Engine Government does the following:

- launches control plane components as static pods, managed by the kubelet. The embedded container runtime is containerd.
- provides defaults and configuration options that allow clusters to pass the CIS Kubernetes Benchmark v1.5 or v1.6 with minimal operator intervention

- enables FIPS 140-2 compliance
- regularly scans components for CVEs using trivy in our build pipeline

With Rancher Kubernetes Engine Government we take lessons learned from developing and maintaining our lightweight Kubernetes distribution, K3s, and apply them to build an enterprise-ready distribution with K3s ease-of-use. What this means is that Rancher Kubernetes Engine Government is, at its simplest, a single binary to be installed and configured on all nodes expected to participate in the Kubernetes cluster. When started, Rancher Kubernetes Engine Government is then able to bootstrap and supervise role-appropriate agents per node while sourcing needed content from the network.

The fundamental roles for the nodes and core functionality of Rancher Kubernetes Engine Government are represented in the following figure:

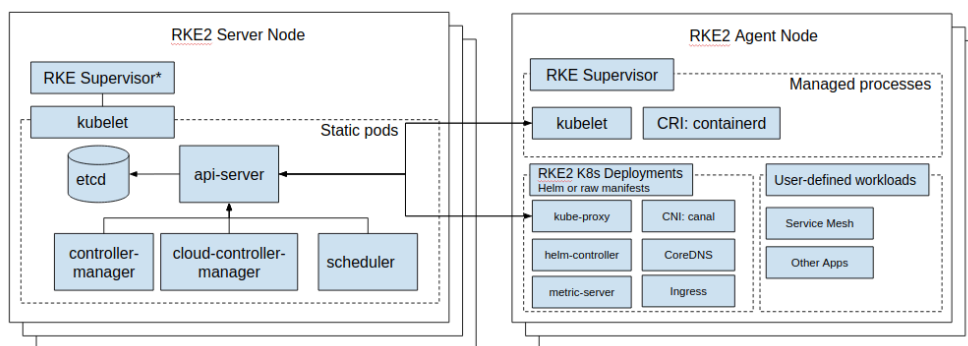


FIGURE 4.1: COMPONENT OVERVIEW - RANCHER KUBERNETES ENGINE GOVERNMENT

Rancher Kubernetes Engine Government brings together several open source technologies to make this all work:

- K3s - [Helm Controller \(https://github.com/k3s-io/helm-controller\)](https://github.com/k3s-io/helm-controller) ↗
- Kubernetes
 - API Server
 - Controller Manager
 - Kubelet
 - Scheduler
 - Proxy
- etcd

- Container Runtime - runc, containerd/cri
- CoreDNS
- NGINX Ingress Controller
- Metrics Server
- Helm

All of these, except the NGINX Ingress Controller, are compiled and statically linked with ¹

While all Rancher Kubernetes Engine Government roles can be installed on a single system, for the best availability, performance and security, the recommended deployment of a Rancher Kubernetes Engine Government cluster is a pair of nodes for the control plane role, at least three etcd role-based nodes and three or more worker nodes.

Rancher Kubernetes Engine Government can run as a complete cluster on a single node or can be expanded into a multi-node cluster. Besides the core Kubernetes components, these are also configurable and included:

- Multiple Kubernetes versions
- CoreDNS, Metrics, Ingress controller
- CNI: Canal (Calico & Flannel), Cilium or Calico
- Fleet Agent : for GitOps deployment of cloud-native applications

All of these components are configurable and can be swapped out for your implementation of choice. With these included components, you get a fully functional and CNCF-conformant cluster so you can start running apps right away.



Tip

Learn more information about Rancher Kubernetes Engine Government at <https://documentation.suse.com/cloudnative/rke2/> [↗](#).

While all Rancher Kubernetes Engine Government roles can be installed on a single system, a multi-node cluster, is a more production-like approach and will be described in the deployment section.

¹ <https://github.com/golang/go/tree/dev.boringcrypto/misc/boring> [↗](#)



Tip

To improve availability, performance and security, the recommended deployment of a Rancher Kubernetes Engine Government cluster is a pair of nodes for the control plane role, at least three etcd role-based nodes and three or more worker nodes.

4.3 Software - SUSE Linux Enterprise Server

SUSE Linux Enterprise Server (SLES (<https://www.suse.com/products/server/>)) is an adaptable and easy-to-manage platform that allows developers and administrators to deploy business-critical workloads on-premises, in the cloud and at the edge. It is a Linux operating system that is adaptable to any environment – optimized for performance, security and reliability. As a multimodal operating system that paves the way for IT transformation in the software-defined era, this simplifies multimodal IT, makes traditional IT infrastructure efficient and provides an engaging platform for developers. As a result, one can easily deploy and transition business-critical workloads across on-premises and public cloud environments.

Designed for interoperability, SUSE Linux Enterprise Server integrates into classical Unix and Windows environments, supports open standard interfaces for systems management, and has been certified for IPv6 compatibility. This modular, general purpose operating system runs on four processor architectures and is available with optional extensions that provide advanced capabilities for tasks such as real time computing and high availability clustering. SUSE Linux Enterprise Server is optimized to run as a high performing guest on leading hypervisors and supports an unlimited number of virtual machines per physical system with a single subscription. This makes it the perfect guest operating system for virtual computing.

4.4 Compute Platform

Leveraging the enterprise grade functionality of the operating system mentioned in the previous section, many compute platforms can be the foundation of the deployment:

- Virtual machines on supported hypervisors or hosted on cloud service providers
- Physical, baremetal or single-board computers, either on-premises or hosted by cloud service providers



Note

To complete self-testing of hardware with [SUSE YES Certified Process \(https://www.suse.com/partners/ihv/yes/yes-certified-process\)](https://www.suse.com/partners/ihv/yes/yes-certified-process), you can download and install the respective SUSE operating system support-pack version of SUSE Linux Enterprise Server and the YES test suite. Then run the tests per the instructions in the test kit, fixing any problems encountered and when corrected, re-run all tests to obtain clean test results. Submit the test results into the SUSE Bulletin System (SBS) for audit, review and validation.



Tip

Certified systems and hypervisors can be verified via [SUSE YES Certified Bulletins \(https://www.suse.com/yessearch/\)](https://www.suse.com/yessearch/) and then can be leveraged as supported nodes for this deployment, as long as the certification refers to the respective version of the underlying SUSE operating system required.

[Dell EMC PowerEdge Rack Servers \(https://i.dell.com/sites/csdocuments/Product_Docs/en/powered-edge-rack-servers-quick-reference-guide.pdf\)](https://i.dell.com/sites/csdocuments/Product_Docs/en/powered-edge-rack-servers-quick-reference-guide.pdf) help you build a modern infrastructure that minimizes IT challenges and drives business success. Choose from a complete portfolio of 1, 2, and 4-socket rack servers to deliver high core density for your traditional applications, virtualization, and cloud-native workloads. Enhanced memory speeds, faster NVMe storage options, and BIOS tuning allows you to match performance to your workload for ultimate efficiency

4.4.1 Dell EMC PowerEdge R640 Rack Servers



Note

The EMC PowerEdge R640 is SUSE YES Certified Hardware.

Uncompromising performance and density

The [Dell EMC PowerEdge R640 \(https://www.dell.com/en-us/work/shop/productdetail-stxn/powered-edge-r640\)](https://www.dell.com/en-us/work/shop/productdetail-stxn/powered-edge-r640) is the ideal dual-socket, 1U platform for dense scale-out data center computing. The R640 combines density, performance and scalability to optimize application performance and data center density

- CPU : Up to two 2nd Generation Intel Xeon Scalable processors with up to 28 cores per processor
- Accelerators : Up to one single-width FPGA, Up to three single-width GPU (NVIDIA T4)
- Memory
 - Speed : DIMM Speed (Up to 2933MT/s)
 - Type : RDIMM, LRDIMM, NVDIMM, DCPMM (Intel Optane) DC persistent memory)
 - Module Slots : 24 DDR4 DIMM slots (12 NVDIMM or 12 DCPMM only) Supports registered ECC DDR4 DIMMs only
 - RAM : RDIMM 1.53TB, LRDIMM 3TB, NVDIMM 192GB, DCPMM 6.14TB (7.68TB with LRDIMM)
- Storage
 - Front Bays : Up to 10 x 2.5" with up to 8 NVMe, SAS/SATA/SSD/NVMe, max 76.8TB
 - Up to 10 NVMe, max 64TB
 - Up to 4 x 3.5" SAS/SATA, max 56T
 - Rear Bays : Up to 2 x 2.5" SAS/SATA/SSD/NVMe, max 15.36TB



Tip

For more information, see [Dell EMC PowerEdge R640 Datasheet \(https://i.dell.com/sites/csdocuments/Product_Docs/en/poweredge-r640-spec-sheet.pdf\)](https://i.dell.com/sites/csdocuments/Product_Docs/en/poweredge-r640-spec-sheet.pdf) 

Maximize application performance and density

The scalable business architecture of the R640 is designed to maximize application performance and provide the flexibility to optimize configurations based on the application and use case. With the R640 you can create an NVMe cache pool and use either 2.5" or 3.5" drives for data storage. Combined with up to 24 DIMM's, 12 of which can be DCPMMs or NVDIMMs, you have the resources to maximize application performance with the optimum configuration in only a 1U chassis.

- Simplify deployments and speed deployments with Dell EMC ready nodes for ScaleIO and VSAN.
- Maximize storage performance with up to 10 NVMe drives or 12 2.5” drives.
- Scale compute resources with 2nd Generation Intel Xeon Scalable processors and tailor performance based on your unique workload requirements.

Automate maintenance with Dell EMC OpenManage

The Dell EMC OpenManage portfolio helps deliver peak efficiency for EMC PowerEdge servers, delivering intelligent, automated management of routine tasks. Combined with unique agent-free management capabilities, the R640 is simply managed, freeing up time for high profile projects.

- Simplify management with the OpenManage Enterprise console, with customized reporting and automatic discovery.
- Take advantage of QuickSync 2 capabilities and gain access to your servers easily through your phone or tablet.

Guard your data center with built in security

Every EMC PowerEdge server is designed as part of a cyber resilient architecture, integrating security into the full server life cycle. The R640 leverages new security features built-into every new EMC PowerEdge server strengthening protection so you can reliably and securely deliver accurate data to your customers no matter where they are. By considering each aspect of system security, from design to retirement, Dell EMC ensures trust and delivers a worry-free, secure infrastructure without compromise.

- Rely on a secure component supply chain to ensure protection from factory to the data center.
- Maintain data safety with cryptographically signed firmware packages and Secure Boot.
- Prevent unauthorized or malicious change with Server Lockdown.
- Wipe all data from storage media including hard drives, SSDs and system memory quickly and securely with System Erase.

4.4.2 Dell EMC PowerEdge R740 Rack Servers



Note

The EMC PowerEdge R740 is SUSE YES Certified Hardware.

Optimized for workload acceleration

The Dell EMC PowerEdge R740 (<https://www.dell.com/en-us/work/shop/povw/powerededge-r740>) is designed to accelerate application performance leveraging accelerator card and storage scalability. The 2-socket, 2U platform has the optimum balance of resources to power the most demanding environments

- CPU : Up to two 2nd Generation Intel Xeon Scalable processors with up to 28 cores per processor
- Accelerator : Up to three 300W or six 150W GPUs, Up to three double-width or four single-width FPGAs
- Memory
 - Type : RDIMM, LRDIMM, NVDIMM, DCPMM (Intel Optane DC persistent memory)
 - Module Slots : 24 DDR4 DIMM slots (12 NVDIMM or 12 DCPMM only), Supports registered ECC DDR4 DIMMs only
 - RAM : RDIMM 1.53TB, LRDIMM 3TB, NVDIMM 192GB, DCPMM 6.14TB (7.68TB with LRDIMM)
- Storage : Front Bays: Up to 16 x 2.5" SAS/SATA/SSD, max 122.88TB, Up to 8 x 3.5" SAS/SATA, max 128TB



Tip

For more information, see Dell EMC PowerEdge R740 Rack Servers (<https://www.dell.com/en-us/work/shop/povw/powerededge-r740>)

Expand and optimize application performance

The scalable business architecture of the R740 can scale up to three 300W or six 150W GPUs, or up to three double-width or four single-width FPGAs. With up to 16 2.5” drives or 8 3.5” drives the R740 provides the versatility to adapt to virtually any application and provides the perfect platform for VDI deployments.

- Scale your VDI deployments with 3 double-width GPUs, supporting up to 50% more users when compared to R730.
- Free up storage space using internal M.2 SSDs optimized for boot.
- Scale compute resources with 2nd Generation Intel Xeon Scalable processors and tailor performance based on your unique workload requirements.

Automate systems management with OpenManage

The Dell EMC OpenManage portfolio helps deliver peak efficiency for EMC PowerEdge servers, delivering intelligent, automated management of routine tasks. Combined with unique agent-free management capabilities, the R740 is simply managed, freeing up time for high profile projects.

- Simplify management with the New OpenManage Enterprise console, with customized reporting and automatic discovery.
- Take advantage of QuickSync 2 capabilities and gain access to your servers easily through your phone or tablet.

Rely on EMC PowerEdge with built-in security

Every EMC PowerEdge server is designed as part of a cyber resilient architecture, integrating security into the full server lifecycle. The R740 leverages new security features built-into every new EMC PowerEdge server strengthening protection so you can reliably and securely deliver accurate data to your customers no matter where they are. By considering each aspect of system security, from design to retirement, Dell EMC ensures trust and delivers a worry-free, secure infrastructure without compromise.

- Rely on a secure component supply chain to ensure protection from factory to the data center.
- Maintain data safety with cryptographically signed firmware packages and Secure Boot.

- Prevent unauthorized or malicious change with Server Lockdown.
- Wipe all data from storage media including hard drives, SSDs and system memory quickly and securely with System Erase

4.4.3 Dell EMC PowerEdge R650 Rack Servers



Note

The EMC PowerEdge R650 is SUSE YES Certified Hardware.

Compelling performance, high scalability, and density

The Dell EMC PowerEdge R650 (https://www.dell.com/en-us/work/shop/cty/pdp/spd/powerededge-r650/pe_r650_14796_vi_vp) is a full-featured enterprise server, designed to optimize workloads performance and data center density

- CPU : Up to two 3rd Generation Intel Xeon Scalable processors, with up to 40 cores per processor
- Memory : 32 DDR4 DIMM slots, supports RDIMM 2 TB max or LRDIMM 4 TB max, speeds up to 3200 MT/s, Up to 16 Intel Persistent Memory 200 series (BPS) slots, 8 TB max, Supports registered ECC DDR4 DIMMs only
- Storage Controllers
 - Internal controllers: PERC H745, HBA355I, S150, H345, H755, H755N
 - Boot Optimized Storage Subsystem (BOSS-S2): HW RAID 2 x M.2 SSDs 240 GB or 480 GB
 - External PERC (RAID): PERC H840, HBA355E
- Drive Bays
 - Front bays : Up to 10 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 153 TB, Up to 4 x 3.5-inch SAS/SATA (HDD/SSD) max 64 TB, Up to 8 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 122.8 TB
 - Rear bays : Up to 2 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 30.7 TB



Tip

For more information, see [Dell EMC PowerEdge R650 Datasheet \(https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/dell-emc-powerededge-r650-spec-sheet.pdf\)](https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/dell-emc-powerededge-r650-spec-sheet.pdf) ↗

Innovate at scale with challenging and emerging workloads

The Dell EMC PowerEdge R650, powered by the 3rd Generation Intel Xeon Scalable processors is the optimal rack server to address application performance and acceleration. The EMC PowerEdge R650, is a dual-socket/1U rack server that delivers outstanding performance for the most demanding workloads. It supports 8 channels of memory per CPU, and up to 32 DDR4 DIMMs @ 3200 MT/s speeds. In addition, to address substantial throughput improvements the EMC PowerEdge R650 supports PCIe Gen 4 and up to 10 NVMe drives with improved air-cooling features and optional Direct Liquid Cooling to support increasing power and thermal requirements. This makes the EMC PowerEdge R650 an ideal server for data center standardization on a wide range of workloads including; Database and Analytics, HighFrequency Trading, Traditional corporate IT, Virtual Desktop Infrastructure, and even HPC or AI/ML environments that require performance, and GPU support in a dense 1U form factor

Increase efficiency and accelerate operations with autonomous collaboration

The Dell EMC OpenManage systems management portfolio tames the complexity of managing and securing IT infrastructure. Using Dell Technologies' intuitive end-to-end tools, IT can deliver a secure, integrated experience by reducing process and information silos in order to focus on growing the business. The Dell EMC OpenManage portfolio is the key to your innovation engine, unlocking the tools and automation that help you scale, manage, and protect your technology environment.

- Built-in telemetry streaming, thermal management, and RESTful API with Redfish offer streamlined visibility and control for better server management
- Intelligent automation lets you enable cooperation between human actions and system capabilities for added productivity
- Integrated change management capabilities for update planning and seamless, zero-touch configuration and implementation
- Full-stack management integration with Microsoft, VMware, ServiceNow, Ansible and many other tools

Protect your data assets and infrastructure with proactive resilience

The Dell EMC PowerEdge R650 server is designed with a cyber-resilient architecture, integrating security deeply into every phase in the lifecycle, from design to retirement.

- Operate your workloads on a secure platform anchored by cryptographically trusted booting and silicon root of trust
- Maintain server firmware safety with digitally signed firmware packages
- Prevent unauthorized configuration or firmware change with system lockdown
- Securely and quickly wipe all data from storage media, including hard drives, SSDs and system memory with System Erase

4.4.4 Dell EMC PowerEdge R750 Rack Servers



Note

The EMC PowerEdge R750 is SUSE YES Certified Hardware.

General purpose server optimized to address the most demanding workloads

The Dell EMC PowerEdge R750 (https://www.dell.com/en-us/work/shop/cty/pdp/spd/powerededge-r750/pe_r750_14794_vi_vp) is a full-featured enterprise server, delivering outstanding performance for the most demanding workloads

- CPU : Up to two 3rd Generation Intel Xeon Scalable processors, with up to 40 cores per processor
- Memory : 32 DDR4 DIMM slots, supports RDIMM 2 TB max or LRDIMM 8 TB max, speeds up to 3200 MT/s, Up to 16 Intel Persistent Memory 200 series (BPS) slots, 8 TB max, Supports registered ECC DDR4 DIMMs only
- Storage Controllers

- Internal controllers : PERC H745, HBA355I, S150, H345, H755, H755N
- Boot Optimized Storage Subsystem (BOSS-S2): HW RAID 2 x M.2 SSDs 240 GB or 480 GB
- External PERC (RAID): PERC H840, HBA355E
- Drive Bays
 - Front bays : Up to 12 x 3.5-inch SAS/SATA (HDD/SSD) max 192 TB, Up to 8 x 2.5-inch NVMe (SSD) max 122.88 TB, Up to 16 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 245.76 TB, Up to 24 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 368.84 TB
 - Rear bays : Rear bays: Up to 2 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 30.72 TB, Up to 4 x 2.5-inch SAS/SATA/NVMe (HDD/SSD) max 61.44 TB



Tip

For more information, see see [Dell EMC PowerEdge R750 Datasheet](https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/dell-emc-powerededge-r750-spec-sheet.pdf) (<https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/dell-emc-powerededge-r750-spec-sheet.pdf>) ↗

Innovate at scale with challenging and emerging workloads

The Dell EMC PowerEdge R750, powered by the 3rd Generation Intel Xeon Scalable processors is a rack server to address application performance and acceleration. The EMC PowerEdge R750, is a dual-socket/2U rack server that delivers outstanding performance for the most demanding workloads. It supports 8 channels of memory per CPU, and up to 32 DDR4 DIMMs @ 3200 MT/s speeds. In addition, to address substantial throughput improvements the EMC PowerEdge R750 supports PCIe Gen 4 and up to 24 NVMe drives with improved air-cooling features and optional Direct Liquid Cooling to support increasing power and thermal requirements. This makes the EMC PowerEdge R750 an ideal server for data center standardization on a wide range of workloads including; Database and Analytics, High Performance Computing (HPC), Traditional corporate IT, Virtual Desktop Infrastructure, and AI/ML environments that require performance, extensive storage and GPU support

Increase efficiency and accelerate operations with autonomous collaboration

The Dell EMC OpenManage systems management portfolio tames the complexity of managing and securing IT infrastructure. Using Dell Technologies' intuitive end-to-end tools, IT can deliver a secure, integrated experience by reducing process and information silos in order to focus on growing the business. The Dell EMC OpenManage portfolio is the key to your innovation engine, unlocking the tools and automation that help you scale, manage, and protect your technology environment.

- Built-in telemetry streaming, thermal management, and RESTful API with Redfish offer streamlined visibility and control for better server management
- Intelligent automation lets you enable cooperation between human actions and system capabilities for added productivity
- Integrated change management capabilities for update planning and seamless, zero-touch configuration and implementation
- Full-stack management integration with Microsoft, VMware, ServiceNow, Ansible and many other tool

Protect your data assets and infrastructure with proactive resilience

The Dell EMC PowerEdge R750 server is designed with a cyber-resilient architecture, integrating security deeply into every phase in the lifecycle, from design to retirement.

- Operate your workloads on a secure platform anchored by cryptographically trusted booting and silicon root of trust
- Maintain server firmware safety with digitally signed firmware packages
- Prevent unauthorized configuration or firmware change with system lockdown
- Securely and quickly wipe all data from storage media, including hard drives, SSDs and system memory with System Erase



Note

A sample bill of materials, in the [Chapter 9, Appendix](#), cites the necessary quantities of all components, along with a reference to the minimum resource requirements needed by the software components.

5 Deployment

This section describes the process steps for the deployment of the Rancher Kubernetes Engine Government solution. It describes the process steps to deploy each of the component layers starting as a base functional *proof-of-concept*, having considerations on migration toward *production*, providing *scaling* guidance that is needed to create the solution.

5.1 Deployment overview

The deployment stack is represented in the following figure:

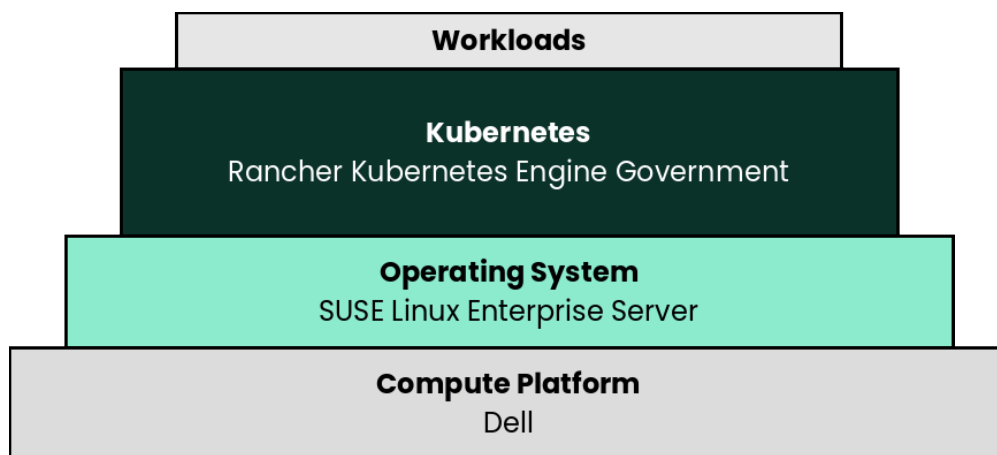


FIGURE 5.1: DEPLOYMENT STACK - RANCHER KUBERNETES ENGINE GOVERNMENT

and details are covered for each layer in the following sections.





Note

The following section's content is ordered and described from the bottom layer up to the top.

5.2 Compute Platform

The base, starting configuration can reside all within a single server. Based upon the relatively small resource requirements for a Rancher Kubernetes Engine Government deployment, a viable approach is to deploy as a virtual machine (VM) on the target nodes, on top of an existing hypervisor, like KVM. For physical host, there are tools that can be used during the setup of the server, see below.

The Integrated Dell Remote Access Controller (iDRAC (<https://www.dell.com/en-ca/lp/dt/open-manage-idrac>) ) is designed for secure local and remote server management and helps IT administrators deploy, update and monitor Dell EMC PowerEdge servers anywhere, anytime.

The Virtual Media (<https://dell.com/support/kbdoc/en-ca/000124001/using-the-virtual-media-function-on-idrac-6-7-8-and-9>) ) function allows the remote usage of software image files (ISO files), which can be used for installing operating systems or updating servers.

5.3 SUSE Linux Enterprise Server

As the base software layer, use an enterprise-grade Linux operating system. For example, SUSE Linux Enterprise Server.

Preparation(s)

To meet the solution stack prerequisites and requirements, SUSE operating system offerings, like [SUSE Linux Enterprise Server](https://www.suse.com/products/server/) (<https://www.suse.com/products/server/>) ) can be used.

1. Ensure these services are in place and configured for this node to use:

- Domain Name Service (DNS) - an external network-accessible service to map IP Addresses to host names
- Network Time Protocol (NTP) - an external network-accessible service to obtain and synchronize system times to aid in time stamp consistency
- Software Update Service - access to a network-based repository for software update packages. This can be accessed directly from each node via registration to

- the general, internet-based [SUSE Customer Center \(https://scc.suse.com\)](https://scc.suse.com) (SCC) or
- an organization's [SUSE Manager \(https://www.suse.com/products/suse-manager/\)](https://www.suse.com/products/suse-manager/) infrastructure or
- a local server running an instance of [Repository Mirroring Tool \(https://documentation.suse.com/sles/15-SP3/single-html/SLES-rmt/#book-rmt\)](https://documentation.suse.com/sles/15-SP3/single-html/SLES-rmt/#book-rmt) (RMT)



Note

During the node's installation, it can be pointed to the respective update service. This can also be accomplished post-installation with the command line tool named [SUSEConnect \(https://www.suse.com/support/kb/doc/?id=000018564\)](https://www.suse.com/support/kb/doc/?id=000018564).

Deployment Process

On the compute platform node, install the noted SUSE operating system, by following these steps:

1. Download the [SUSE Linux Enterprise Server \(https://www.suse.com/download/sles/\)](https://www.suse.com/download/sles/) product (either for the ISO or Virtual Machine image)
 - Identify the appropriate, supported version of SUSE Linux Enterprise Server by reviewing the support matrix for [SUSE Rancher \(https://www.suse.com/suse-rancher/support-matrix/all-supported-versions/\)](https://www.suse.com/suse-rancher/support-matrix/all-supported-versions/) versions Web page.
2. The installation process is described and can be performed with default values by following steps from the product documentation, see [Installation Quick Start \(https://documentation.suse.com/sles/15-SP3/single-html/SLES-installation/#article-installation\)](https://documentation.suse.com/sles/15-SP3/single-html/SLES-installation/#article-installation)



Tip

Adjust both the password and the local network addressing setup to comply with local environment guidelines and requirements.

Deployment Consideration(s)

To further optimize deployment factors, leverage the following practices:

- *Automation*
 - To reduce user intervention, unattended deployments of SUSE Linux Enterprise Server can be automated
 - for ISO-based installations, by referring to the [AutoYaST Guide \(https://documentation.suse.com/sles/15-SP3/single-html/SLES-autoyast/#book-autoyast\)](https://documentation.suse.com/sles/15-SP3/single-html/SLES-autoyast/#book-autoyast) ↗

5.4 Rancher Kubernetes Engine Government

Preparation(s)

1. Identify the appropriate, desired version of the Rancher Kubernetes Engine Government (for example vX.YY.ZZ + rke2rV) by reviewing
 - the "Supported Rancher Kubernetes Engine Government Versions" associated with the respective [SUSE Rancher \(https://www.suse.com/suse-rancher/support-matrix/all-supported-versions/\)](https://www.suse.com/suse-rancher/support-matrix/all-supported-versions/) ↗ version from "Rancher Kubernetes Engine Government Downstream Clusters" section, or
 - the "Releases" on the [Download \(https://github.com/rancher/rke2/\)](https://github.com/rancher/rke2/) ↗ Web page.
 2. For Rancher Kubernetes Engine Government versions 1.21 and higher, if the host kernel supports AppArmor, the AppArmor tools (usually available via the "apparmor-parser" package) must also be present prior to installing Rancher Kubernetes Engine Government.
 - On the SUSE Linux Enterprise Server node, install this required package
- ```
zypper install apparmor-parser
```

3. For the underlying operating system firewall service, either
  - enable and configure the necessary inbound [ports \(https://rancher.com/docs/rke/latest/en/os/#ports\)](https://rancher.com/docs/rke/latest/en/os/#ports) ↗ or
  - stop and completely disable the firewall service.

## Deployment Process

Perform the following steps to install the first Rancher Kubernetes Engine Government server on one of the nodes to be used for the Kubernetes control plane

1. Set the following variable with the noted version of Rancher Kubernetes Engine Government, as found during the preparation steps.

```
RKE2_VERSION=""
```

2. Install the appropriate version of Rancher Kubernetes Engine Government:

- Download the installer script:

```
curl -sL https://get.rke2.io | \
INSTALL_RKE2_VERSION=${RKE2_VERSION} sh -
```

- Set the following variable with the URL that will be used to access the SUSE Rancher server. This may be based on one or more DNS entries, a reverse-proxy server, or a load balancer:

```
RKE2_subjectAltName=
```

- Create the RKE2 config.yaml file:

```
mkdir -p /etc/rancher/rke2/
cat <<EOF> /etc/rancher/rke2/config.yaml
write-kubeconfig-mode: "0644"
tls-san:
 - "${RKE2_subjectAltName}"
EOF
```

3. Start and enable the RKE2 service, which will begin installing the required Kubernetes components:

```
systemctl enable --now rke2-server.service
```

- Include the Rancher Kubernetes Engine Government binary directories in this user's path:

```
echo "PATH=${PATH}:/opt/rke2/bin:/var/lib/rancher/rke2/bin/" >> ~/.bashrc
```

```
source ~/.bashrc
```

- Monitor the progress of the installation:

```
export KUBECONFIG=/etc/rancher/rke2/rke2.yaml
watch -c "kubectl get deployments -A"
```



## Note

For the first two to three minutes of the installation, the initial output will include the error phrase "The connection to the server 127.0.0.1:6443 was refused - did you specify the right host or port?". As Kubernetes services get started this will be replaced with "No resources found". About four minutes after beginning the installation, the output will begin showing the deployments being created, and after six to seven minutes the installation should be complete.

- The Rancher Kubernetes Engine Government deployment is complete when elements of all the deployments (coredns, ingress, and metrics-server) show at least "1" as "AVAILABLE"
  - Use Ctrl+c to exit the watch loop after all deployment pods are running

## Deployment Consideration(s)

To further optimize deployment factors, leverage the following practices:

- *Availability*
  - A full high-availability Rancher Kubernetes Engine Government cluster is recommended for production workloads. The etcd key/value store (aka database) requires an odd number of servers (aka master nodes) be allocated to the Rancher Kubernetes Engine Government cluster. In this case, two additional control-plane servers should be added; for a total of three.



1. Deploy the same operating system on the new compute platform nodes
2. Log in to the first server node and create a new config.yaml file for the remaining two server nodes:

- Set the following variables, as appropriate for this cluster

```
Private IP preferred, if available
FIRST_SERVER_IP=""

Private IP preferred, if available
SECOND_SERVER_IP=""

Private IP preferred, if available
THIRD_SERVER_IP=""

From the /var/lib/rancher/rke2/server/node-token file on the
first server
NODE_TOKEN=""

Match the first of the first server (Hint: `kubectl get
nodes`)
RKE2_VERSION=""
```

- Create the new config.yaml file:

```
echo "server: https://${FIRST_SERVER_IP}:9345" > config.yaml
echo "token: ${NODE_TOKEN}" >> config.yaml
cat /etc/rancher/rke2/config.yaml >> config.yaml
```



## Tip

The next steps require using SCP and SSH. Setting up passwordless SSH, and/or using [ssh-agent](#), from the first server node to the second and third nodes will make these steps quicker and easier.

- Copy the new config.yaml file to the remaining two server nodes:

```
scp config.yaml ${SECOND_SERVER_IP}:~/
scp config.yaml ${THIRD_SERVER_IP}:~/
```

- Move the config.yaml file to the correct location in the file system:

```
ssh ${SECOND_SERVER_IP} << EOF
mkdir -p /etc/rancher/rke2/
cp ~/config.yaml /etc/rancher/rke2/config.yaml
cat /etc/rancher/rke2/config.yaml
EOF

ssh ${THIRD_SERVER_IP} << EOF
mkdir -p /etc/rancher/rke2/
cp ~/config.yaml /etc/rancher/rke2/config.yaml
cat /etc/rancher/rke2/config.yaml
EOF
```

- Execute the following sets of commands on each of the remaining control-plane nodes:

- Install Rancher Kubernetes Engine Government

```
ssh ${SECOND_SERVER_IP} << EOF
curl -sL https://get.rke2.io | \
 INSTALL_RKE2_VERSION=${RKE2_VERSION} sh -
systemctl enable --now rke2-server.service
EOF

ssh ${THIRD_SERVER_IP} << EOF
curl -sL https://get.rke2.io | \
 INSTALL_RKE2_VERSION=${RKE2_VERSION} sh -
systemctl enable --now rke2-server.service
EOF
```

- Monitor the progress of the new server nodes joining the Rancher Kubernetes Engine Government cluster: watch -c "kubectl get nodes"

- It takes up to eight minutes for each node to join the cluster
- A node has deployed correctly when its status is "Ready" and it holds the roles of "control-plane,etcd,master"
- Use Ctrl + c to exit the watch loop after all deployment pods are running



## Note

This can be changed to the normal Kubernetes default by adding a taint to each server node. See the official Kubernetes documentation for more information on how to do that.

3. (Optional) In cases where agent nodes are desired, execute the following sets of commands, using the same, "*RKE2\_VERSION*", "*FIRST\_SERVER\_IP*" and "*NODE\_TOKEN*" variable settings as above, on each of the agent nodes to add it to the Rancher Kubernetes Engine Government cluster:

```
curl -sfL https://get.rke2.io | \
INSTALL_RKE2_VERSION=${RKE2_VERSION} \
RKE2_URL=https://${FIRST_SERVER_IP}:6443 \
RKE2_TOKEN=${NODE_TOKEN} \
RKE2_KUBECONFIG_MODE="644" \
sh -
```

After this successful deployment of the Rancher Kubernetes Engine Government solution, review the [product documentation \(https://documentation.suse.com/cloudnative/rke2/\)](https://documentation.suse.com/cloudnative/rke2/) for details on how to directly use this Kubernetes cluster. Furthermore, by reviewing the SUSE Rancher [product documentation \(https://documentation.suse.com/cloudnative/rancher-manager/\)](https://documentation.suse.com/cloudnative/rancher-manager/) this solution can also be:

- imported (refer to sub-section "Importing Existing Clusters"), then
- managed (refer to sub-section "Cluster Administration") and
- accessed (refer to sub-section "Cluster Access") to address orchestration of workloads, maintaining security and many more functions are readily available.

## 6 Summary

Using components and offerings from [SUSE \(https://www.suse.com\)](https://www.suse.com) and the Rancher portfolio plus [Dell EMC PowerEdge \(https://i.dell.com/sites/csdocuments/Product\\_Docs/en/powerededge-rack-servers-quick-reference-guide.pdf\)](https://i.dell.com/sites/csdocuments/Product_Docs/en/powerededge-rack-servers-quick-reference-guide.pdf) Rack Servers streamline the ability to quickly and effectively engage in a digital transformation, taking advantage of cloud-native resources and disciplines. Using such technology approaches lets you deploy and leverage transformations of infrastructure into a durable, reliable enterprise-grade environment.

### Simplify

Simplify and optimize your existing IT environments

- Using Rancher Kubernetes Engine Government enables you to simplify, maintain and scale Kubernetes cluster deployments in a supportable fashion, with a primary focus on security aspects as well.

### Modernize

Bring applications and data into modern computing

- With Rancher Kubernetes Engine Government, the digital transformation to containerized applications can benefit from the provided, production-quality application infrastructures for each of the respective user bases and to facilitate the actual workload deployments and resilient usage.



### Accelerate

Accelerate business transformation through the power of open source software

- Given the open source nature of Rancher Kubernetes Engine Government and the underlying software components, you can simplify deployment with automation, maintain secure production instance and make significant IT savings as you scale orchestrated microservice deployments anywhere you need to and for whatever use cases are needed, in an agile and innovative way.

## 7 References



### WHITE PAPERS

- **A Buyer's Guide to Enterprise Kubernetes Management Platforms** - [https://more.suse.com/FY22\\_Buyers\\_Guide\\_to\\_Enterprise\\_Container\\_Management\\_Buyers-Guide-to-Kubernetes-Management-Platforms.html](https://more.suse.com/FY22_Buyers_Guide_to_Enterprise_Container_Management_Buyers-Guide-to-Kubernetes-Management-Platforms.html) 
- **How to Build an Enterprise Kubernetes Strategy** - <https://more.suse.com/FY22-global-web-How-to-Build-Enterprise-K8s-Strategy.html> 













### BOOKS

- **Kubernetes Management** - <https://more.suse.com/rs/937-DCH-261/images/002022021-DummiesGuide.pdf> 

### TRAINING

- **SUSE** - <https://training.suse.com/> 
- **Rancher** - <https://rancher.com/training/> 

### WEB SITES

- **SUSE** - <https://www.suse.com> 
- **SUSE Customer Center (SCC)** - <https://scc.suse.com> 
- **Products**
  - **SUSE Rancher** - <https://www.suse.com/products/rancher/>  (documentation (<https://documentation.suse.com/cloudnative/rancher-manager/>) 
  - **Rancher Kubernetes Engine (RKE)** - <https://rancher.com/products/rke/>  (documentation (<https://rancher.com/docs/rke/latest/en/>) 
  - **K3s** - <https://www.suse.com/products/k3s/>  (documentation (<https://documentation.suse.com/cloudnative/k3s/>) 
  - **SUSE Linux Enterprise Micro (SLEMicro)** - <https://www.suse.com/products/micro/>  (documentation (<https://documentation.suse.com/sle-micro/5.5/>) 
  - **SUSE Linux Enterprise Server (SLES)** - <https://www.suse.com/products/server/>  (documentation (<https://documentation.suse.com/sles/15-SP3/>) 

- SUSE Manager - <https://www.suse.com/products/suse-manager/> (documentation (<https://documentation.suse.com/suma/4.3/>))
- SUSE Repository Mirroring Tool (RMT) - <https://www.suse.com/products/server/> (documentation (<https://documentation.suse.com/sles/15-SP3/single-html/SLES-rmt/#book-rmt>))
- Projects
  - Rancher Kubernetes Engine Government (RKE2) - <https://github.com/rancher/rke2> (documentation (<https://docs.rke2.io/>))
- Dell - <https://www.dell.com/en-us>
  - Dell EMC PowerEdge Rack Servers - [https://i.dell.com/sites/csdocuments/Product\\_Docs/en/poweredge-rack-servers-quick-reference-guide.pdf](https://i.dell.com/sites/csdocuments/Product_Docs/en/poweredge-rack-servers-quick-reference-guide.pdf)

## 8 Glossary

- Document Scope

### Reference Configuration

A guide with the basic steps to deploy the layered stack of components from both the SUSE and partner portfolios. This is considered a fundamental basis to demonstrate a specific, tested configuration of components.

### Reference Architectures<sup>1</sup>

A guide with the general steps to deploy and validate the structured solution components from both the SUSE and partner portfolios. This provides a shareable template of consistency for consumers to leverage for similar production ready solutions, including design considerations, implementation suggestions and best practices.

### Best Practice

Information that can overlap both the SUSE and partner space. It can either be provided as a stand-alone guide that provides reliable technical information not covered in other product documentation, based on real-life installation and implementation experiences from subject matter experts or complementary, embedded sections within any of the above documentation types describing considerations and possible steps forward.

- Factor(s)

### Automation<sup>2</sup>

Infrastructure automation enables speed through faster execution when configuring the infrastructure and aims at providing visibility to help other teams across the enterprise work quickly and more efficiently. Automation removes the risk associated with human error, like manual misconfiguration; removing this can decrease downtime and increase reliability. These outcomes and attributes help the enterprise move toward implementing a culture of DevOps, the combined working of development and operations.

---

<sup>1</sup> link: [Reference Architecture \(https://en.wikipedia.org/wiki/Reference\\_architecture\)](https://en.wikipedia.org/wiki/Reference_architecture) ↗

<sup>2</sup> link: [Infrastructure-as-Code \(https://en.wikipedia.org/wiki/Infrastructure\\_as\\_code\)](https://en.wikipedia.org/wiki/Infrastructure_as_code) ↗

### Availability<sup>3</sup>

The probability that an item operates satisfactorily, without failure or downtime, under stated conditions as a function of its reliability, redundancy and maintainability attributes. Some major objectives to achieve a desired service level objectives are:

- Preventing or reducing the likelihood and frequency of failures via design decisions within the allowed cost of ownership
- Correcting or coping with possible component failures via resiliency, automated failover and disaster-recovery processes
- Estimating and analyzing current conditions to prevent unexpected failures via predictive maintenance

### Integrity<sup>4</sup>

Integrity is the maintenance of, and the insurance of the accuracy and consistency of a specific element over its entire lifecycle. Both physical and logical aspects must be managed to ensure stability, performance, re-usability and maintainability.

### Security<sup>5</sup>

Security is about ensuring freedom from or resilience against potential harm, including protection from destructive or hostile forces. To minimize risks, one must manage governance to avoid tampering, maintain access controls to prevent unauthorized usage and integrate layers of defense, reporting and recovery tactics.

- Deployment Flavor(s)

### Proof-of-Concept<sup>6</sup>

A partial or nearly complete prototype constructed to demonstrate functionality and feasibility for verifying specific aspects or concepts under consideration. This is often a starting point when evaluating a new, transitional technology. Sometimes it starts as a Minimum Viable Product (MVP<sup>7</sup>) that has just enough features to satisfy an

---

3 link: [Availability](https://en.wikipedia.org/wiki/Availability) (https://en.wikipedia.org/wiki/Availability) ↗

4 link: [Data Integrity](https://en.wikipedia.org/wiki/Data_integrity) (https://en.wikipedia.org/wiki/Data\_integrity) ↗

5 link: [Security](https://en.wikipedia.org/wiki/Security) (https://en.wikipedia.org/wiki/Security) ↗

6 link: [Proof of Concept](https://en.wikipedia.org/wiki/Proof_of_concept) (https://en.wikipedia.org/wiki/Proof\_of\_concept) ↗

7 link: [Minimum Viable Product](https://en.wikipedia.org/wiki/Minimum_viable_product) (https://en.wikipedia.org/wiki/Minimum\_viable\_product) ↗



initial set of requests. After such insights and feedback are obtained and potentially addressed, redeployments may be used to iteratively branch into other realms or to incorporate other known working functionality.

### **Production**

A deployed environment that target customers or users can interact with and rely upon to meet their needs, plus be operationally sustainable in terms of resource usage and economic constraints.

### **Scaling**

The flexibility of a system environment to either vertically scale-up, horizontally scale-out or conversely scale-down by adding or subtracting resources as needed. Attributes like capacity and performance are often the primary requirements to address, while still maintaining functional consistency and reliability.

## 9 Appendix

The following sections provide a bill of materials listing for the respective component layer(s) of the described deployment.

### 9.1 Compute platform bill of materials

Sample set of computing platform models, components and resources.

| Role      | Qty | SKU      | Component                                                                                                                                | Notes                                                                         |
|-----------|-----|----------|------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------|
| Example 1 | 1-3 |          | <b>Dell EMC PowerEdge R640</b>                                                                                                           | <ul style="list-style-type: none"><li>• items below listed per node</li></ul> |
|           | 1   | 321-BCQQ | <ul style="list-style-type: none"><li>• 2.5 in. chassis with up to 10 hard drives, 8 NVMe drives, and 3 PCIe slots, 2 CPU only</li></ul> |                                                                               |
|           | 2   | 338-BTSI | <ul style="list-style-type: none"><li>• Intel Xeon Gold 6238 2.1 G, 22C/44 T, 10.4 GT/s, 30.25</li></ul>                                 |                                                                               |

| Role | Qty           | SKU      | Component                                                                                                        | Notes |
|------|---------------|----------|------------------------------------------------------------------------------------------------------------------|-------|
|      |               |          | M Cache, Turbo, HT (140 W)<br>DDR4-2933                                                                          |       |
|      | 12            | 370-AEVN | <ul style="list-style-type: none"> <li>32 GB RDIMM, 3200 MT/s, Dual Rank</li> </ul>                              |       |
|      | 1             | 405-AAJU | <ul style="list-style-type: none"> <li>HBA330 12 Gbps SAS HBA Controller (NON-RAID), minicard</li> </ul>         |       |
|      | 1             | 385-BBKT | <ul style="list-style-type: none"> <li>iDRAC9, Enterprise</li> </ul>                                             |       |
|      | 1             | 330-BBGN | <ul style="list-style-type: none"> <li>Riser Config 2, 3 x 16 LP</li> </ul>                                      |       |
|      | 1             | 406-BBLG | <ul style="list-style-type: none"> <li>Mellanox ConnectX-4 Lx Dual Port 25 GbE SFP 28 rNDC</li> </ul>            |       |
|      | 1             | 406-BBLD | <ul style="list-style-type: none"> <li>Mellanox ConnectX-4 Lx dual port 25 GbE SFP28 NIC, low profile</li> </ul> |       |
|      | 1             | 450-ADWS | <ul style="list-style-type: none"> <li>Dual, hot-plug, redundant power supply (1 + 1), 750 W</li> </ul>          |       |
|      | 2 min – 8 max | 400-BELT | <ul style="list-style-type: none"> <li>Dell 1.6 TB, NVMe, Mixed Use Express Flash, 2.5</li> </ul>                |       |

| Role      | Qty | SKU      | Component                                                                                                                              | Notes                                                                         |
|-----------|-----|----------|----------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------|
|           |     |          | SFF Drive, U.2,<br>P4610 with Carrier                                                                                                  |                                                                               |
|           | 2   | 400-AZQO | <ul style="list-style-type: none"> <li>800 GB SSD SAS Mix Use 12 Gbps 512e 2.5 in Hot-plug AG Drive, 3 DWPD, 4380 TBW</li> </ul>       |                                                                               |
|           | 1   | 403-BCHI | <ul style="list-style-type: none"> <li>BOSS Cntrl + 2 M.2 240G, R1, LP1</li> </ul>                                                     |                                                                               |
| Example 2 | 1-3 |          | <b>Dell EMC PowerEdge R740</b>                                                                                                         | <ul style="list-style-type: none"> <li>items below listed per node</li> </ul> |
|           | 1   | 321-BCRC | <ul style="list-style-type: none"> <li>Chassis up to 24 x 2.5 in. hard drives including 12 NVME drives, 2 CPU configuration</li> </ul> |                                                                               |
|           | 1   | 338-BTSI | <ul style="list-style-type: none"> <li>Intel Xeon Gold 6238 2.1 G, 22C/44 T, 10.4 GT/s, 30.25</li> </ul>                               |                                                                               |

| Role | Qty     | SKU      | Component                                                                                                                    | Notes |
|------|---------|----------|------------------------------------------------------------------------------------------------------------------------------|-------|
|      |         |          | M Cache, Turbo, HT (140 W)<br>DDR4-2933                                                                                      |       |
|      | 12      | 370-AEVN | <ul style="list-style-type: none"> <li>32 GB RDIMM, 2933MT/s, Dual Rank</li> </ul>                                           |       |
|      | 1       | 405-AANK | <ul style="list-style-type: none"> <li>HBA330 controller adapter, low profile</li> </ul>                                     |       |
|      | 1       | 385-BBKT | <ul style="list-style-type: none"> <li>iDRAC9, Enterprise</li> </ul>                                                         |       |
|      | 1       | 330-BBHD | <ul style="list-style-type: none"> <li>Riser Config 6, 5 x 8, 3 x1 6 slots</li> </ul>                                        |       |
|      | 1       | 406-BBLG | <ul style="list-style-type: none"> <li>Mellanox ConnectX-4 Lx Dual Port 25 GbE SFP28 rNDC</li> </ul>                         |       |
|      | 1       | 406-BBLE | <ul style="list-style-type: none"> <li>Mellanox ConnectX-4 Lx Dual Port 25 GbE SFP28 network interface controller</li> </ul> |       |
|      | 1       | 450-ADWM | <ul style="list-style-type: none"> <li>Dual, hot-plug, redundant power supply (1 + 1), 1100 W</li> </ul>                     |       |
|      | 1       | 403-BCHP | <ul style="list-style-type: none"> <li>BOSS Cntrl + 2 M.2 240G, R1, FH</li> </ul>                                            |       |
|      | 1 to 24 | 400-AZQO | <ul style="list-style-type: none"> <li>800 GB, 1.92 TB, or 3.84 TB SSD SAS mixed use 12</li> </ul>                           |       |

| Role | Qty     | SKU      | Component                                                                                                              | Notes |
|------|---------|----------|------------------------------------------------------------------------------------------------------------------------|-------|
|      |         |          | Gbps 512e 2.5 in.<br>hot-plug AG drive with carrier,<br>3 DWPD, 4380 TBW, CK                                           |       |
|      | 1 to 12 | 400-BLKD | <ul style="list-style-type: none"> <li>• Dell 1.6 TB, 3.2 TB, or 6.4 TB, NVMe, mixed use express flash, 2.5</li> </ul> |       |

| Role                                          | Qty | SKU      | Component                                                                                                               | Notes                                                                         |
|-----------------------------------------------|-----|----------|-------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------|
|                                               |     |          | SFF drive, U.2,<br>P4610 with carrier, CK                                                                               |                                                                               |
| Example 3                                     | 1-3 |          | <b>Dell EMC PowerEdge R650</b>                                                                                          | <ul style="list-style-type: none"> <li>items below listed per node</li> </ul> |
|                                               | 1   | 321-BGHH | <ul style="list-style-type: none"> <li>2.5" Chassis with up to 8 Hard Drives (SAS/SATA), 3 PCIe Slots, 2 CPU</li> </ul> |                                                                               |
|                                               | 2   | 338-BZXK | <ul style="list-style-type: none"> <li>Intel Xeon Gold 6330 2G, 42M, 205W</li> </ul>                                    |                                                                               |
|                                               | 16  | 370-AEVQ | <ul style="list-style-type: none"> <li>16GB RDIMM, 3200MT/s, Dual Rank</li> </ul>                                       |                                                                               |
|                                               | 1   | 450-AIQZ | <ul style="list-style-type: none"> <li>Dual, Hot-plug, PSU 1 + 1, 1400W, MM</li> </ul>                                  |                                                                               |
|                                               | 1   | 528-CRVW | <ul style="list-style-type: none"> <li>iDRAC9 Datacenter 15G</li> </ul>                                                 |                                                                               |
|                                               | 1   | 340-CUQN | <ul style="list-style-type: none"> <li>R650 Ship 4x3.5, 10x2.5</li> </ul>                                               |                                                                               |
|                                               | 1   | 321-BGHG | <ul style="list-style-type: none"> <li>8x2.5 Chipset NVMe RAID Config, 2CPU</li> </ul>                                  |                                                                               |
|                                               | 1   | 405-AAZE | <ul style="list-style-type: none"> <li>PERC H755N Front</li> </ul>                                                      |                                                                               |
|                                               | 1   | 330-BBRP | <ul style="list-style-type: none"> <li>Riser C0-2, 3x16 LP, HL</li> </ul>                                               |                                                                               |
| Sample set of software, support and services. |     |          |                                                                                                                         |                                                                               |
|                                               | 1   | 403-BCMB | <ul style="list-style-type: none"> <li>BOSS-S2 Cntrl + 2 M.2 480G</li> </ul>                                            |                                                                               |
|                                               | 3   | 400-BLKD | <ul style="list-style-type: none"> <li>1.6TB, NVMe, 2.5</li> </ul>                                                      |                                                                               |
| 43                                            |     |          | Dr, M0, P5600                                                                                                           | Software, Bill of Materials   SUSE Linux Enterprise...                        |
|                                               | 1   | 540-BCOF | <ul style="list-style-type: none"> <li>Mlnx ConX5 DP 10/25Gbe SFP28 QCP3.0</li> </ul>                                   |                                                                               |

| Role                    | Qty | SKU        | Component                                                                                                                                     | Notes                                                                                                             |
|-------------------------|-----|------------|-----------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| Operating System        | 1-3 | 874-006875 | SUSE Linux Enterprise Server, <ul style="list-style-type: none"> <li>• x86_64,</li> <li>• Priority Subscription,</li> <li>• 1 Year</li> </ul> | Configuration: <ul style="list-style-type: none"> <li>• per node (up to 2 sockets, stackable) or 2 VMs</li> </ul> |
| Kubernetes Management   | 1   | R-0001-PS1 | SUSE Rancher, <ul style="list-style-type: none"> <li>• x86-64,</li> <li>• Priority Subscription,</li> <li>• 1 Year</li> </ul>                 | Configuration: <ul style="list-style-type: none"> <li>• per deployed instance</li> </ul>                          |
| Rancher Management      | 2   | R-0004-PS1 | Rancher 10 Nodes <ul style="list-style-type: none"> <li>• x86-64 or aarch64,</li> <li>• Priority Subscription,</li> <li>• 1 Year,</li> </ul>  | Configuration: <ul style="list-style-type: none"> <li>• requires priority server subscription</li> </ul>          |
| Consulting and Training | 1   | R-0001-QSO | Rancher Quick Start, <ul style="list-style-type: none"> <li>• Go Live Services</li> </ul>                                                     |                                                                                                                   |



## Note

For the software components, other support term durations are also available.



## 9.3 Documentation configuration / attributes


This document was built using the following [AsciiDoc](https://github.com/asciidoc/asciidoc) (<https://github.com/asciidoc/asciidoc>) and DocBook Authoring and Publishing Suite ([DAPS](https://github.com/openSUSE/daps) (<https://github.com/openSUSE/daps>)) attributes:

```
Appendix=1 Arch0v=1 Automation=1 Availability=1 BP=1 BPBV=1 CompMod=1 DepConsiderations=1
Deployment=1 FCTR=1 FLVR=1 GFDL=1 Glossary=1 HWComp=1 HWDepCfg=1 IHV-Dell-PowerEdge=1
IHV-Dell=1 Integrity=1 LN=1 PoC=1 Production=1 RA=1 RC=1 References=1 Requirements=1
SWComp=1 SWDepCfg=1 Scaling=1 Security=1 docdate=2022-03-28 env-daps=1 focusRKE2=1
iIHV=1 iK3s=1 iRKE1=1 iRKE2=1 iRMT=1 iRancher=1 iSLEMicro=1 iSLES=1 iSUMa=1 layerSLES=1
```

## 10 Legal Notice

Copyright © 2006–2025 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled "GNU Free Documentation License".

SUSE, the SUSE logo and YaST are registered trademarks of SUSE LLC in the United States and other countries. For SUSE trademarks, see <https://www.suse.com/company/legal/> .

Linux is a registered trademark of Linus Torvalds. All other names or trademarks mentioned in this document may be trademarks or registered trademarks of their respective owners.

Documents published as part of the series SUSE Technical Reference Documentation have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

# 11 GNU Free Documentation License

Copyright © 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

## 0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

## 1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition. The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

## 2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

## 3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

## 4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.

- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

## 5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.



In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

## 6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

## 7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

## 8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all

Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

## 9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

## 10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

## ADDENDUM: How to use this License for your documents

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document  
under the terms of the GNU Free Documentation License, Version 1.2

```
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.
A copy of the license is included in the section entitled "GNU
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “ with... Texts.” line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.