Canonical OpenStack on Lenovo ThinkSystem SR630 V4 and SR650 V4: An Open, Sovereign Private Cloud

Solution Brief

Private cloud has become a strategic layer of enterprise IT. It hosts traditional virtual machines, cloud-native applications and, increasingly, AI workloads. IT teams need the agility of public cloud together with control over data location, security posture, and cost.

OpenStack is the de facto open-source standard for building private and sovereign clouds, and banks, governments, telcos, and enterprises have run it in production for more than a decade. Canonical has contributed to OpenStack since the project began and is one of its largest contributors. It also publishes Ubuntu, the most widely used operating system for OpenStack deployments.

Canonical OpenStack is an upstream-aligned distribution with no forks and no proprietary extensions. Model-driven automation turns OpenStack from a complex integration project into a repeatable, supportable product. Running it on Lenovo ThinkSystem SR630 V4 and SR650 V4 servers gives customers a reference architecture for a production cloud, with lower deployment risk and faster time to value.

Business Challenge

Infrastructure teams face pressure on several fronts at once:

- **Escalating licensing costs**: per-core, subscription-only hypervisor licensing and forced product bundling increase total cost of ownership (TCO) and make budgets unpredictable.
- **Vendor lock-in**: proprietary virtualization stacks restrict architectural choice and make it costly to leave.
- **Digital sovereignty and compliance**: regulations and data-residency rules demand full control over where data lives, who operates the platform, and how it is secured.
- **Operational complexity**: integrating and running compute, network, storage, and observability as a cloud requires deep expertise that many teams lack.
- **Long-term security and lifecycle**: infrastructure must stay patched and compliant for many years, and upgrades must not disrupt workloads.
- **Modern and AI workloads**: a single platform must serve VMs, Kubernetes, and GPU-accelerated workloads side by side.

Solution

Canonical OpenStack on Lenovo ThinkSystem SR630 V4 and SR650 V4 is an open, software-defined private cloud built from a small number of well-integrated layers:

- **Lenovo ThinkSystem SR630 V4 and SR650 V4 servers**: the primary hardware platforms for the solution. The 1U SR630 V4 supports dense, scale-out cloud nodes, while the 2U SR650 V4 provides greater expansion flexibility for storage-rich, accelerator-enabled, and mixed-role configurations. Both platforms use Intel Xeon 6 processors, DDR5 memory, PCIe Gen5, NVMe storage options, and Lenovo XClarity Controller management. Both platforms are Ubuntu certified for Ubuntu 22.04, 24.04 and 26.04 LTS.
- **Ubuntu Server LTS with Ubuntu Pro**: the host OS on every node, with Livepatch rebootless kernel patching and long-term security maintenance.
- **Canonical MAAS and Juju**: MAAS (Medium and Large) discovers, tests, and deploys bare-metal servers, while Juju charms deploy, integrate, scale, and upgrade every cloud service consistently.
- **Canonical OpenStack**: Nova on KVM, Neutron with OVN, Cinder, Glance, Keystone, and Horizon, plus optional object storage, load balancing, DNS, file shares, telemetry, and instance high availability.
- **Canonical Ceph, COS, and Landscape**: scale-out block, object, and file storage; integrated metrics, logs, and alerting; and fleet-wide patching and compliance.

The following figure shows how these layers fit together. Lenovo ThinkSystem SR630 V4 and SR650 V4 servers, the XClarity Controller, and a redundant Ethernet fabric form the physical foundation. Ubuntu Server with Ubuntu Pro runs on every node. MAAS and Juju deploy and operate the OpenStack services above it, while COS and Landscape provide observability, patching and compliance. The Small design leaves out MAAS, COS and Landscape, as described in the Reference Configurations section. Users and DevOps teams consume compute, networking, storage, and security services through standard OpenStack APIs and tools, so virtual machines, Kubernetes clusters, and GPU workloads all run on one shared, open platform.

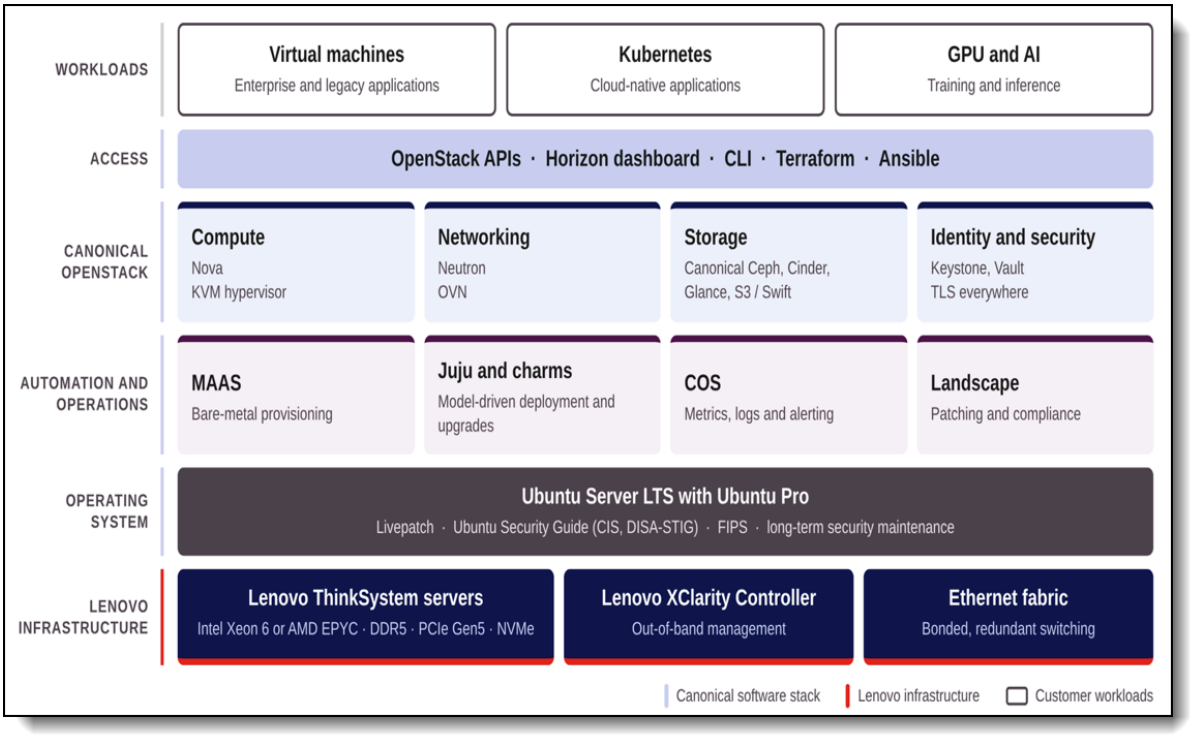


Figure 1. Canonical OpenStack on Lenovo ThinkSystem SR630 V4 and SR650 V4 solution architecture.

Resilient and secure by design

Every control-plane service runs as a three-member cluster spread across availability zones, Ceph keeps three copies of the data, and server networking is bonded to redundant switches, so the design has no single point of failure. Security features include TLS on all endpoints, Vault Secrets management, encryption at rest, DISA-STIG and CIS hardening through the Ubuntu Security Guide, and FIPS certified cryptographic modules for the Ubuntu host. Updates are applied in a rolling fashion to minimize workload disruption, and LTS releases follow a supported upgrade path.

Simplified deployment and management

MAAS and Lenovo XClarity work side by side. MAAS automates the server lifecycle and uses the XClarity Controller for power control over IPMI, while XClarity manages firmware compliance, hardware health and alerts, and remote console. Canonical engineers can design and deploy the cloud on the customer's Lenovo ThinkSystem SR630 V4 or SR650 V4 hardware; customers then operate it themselves with Ubuntu Pro + Support or choose a fully managed service from Canonical, backed by Lenovo services and Premier Support for the hardware.

XCC3 ships with IPMI over LAN turned off, so every medium and large node is ordered with the factory option to enable it (feature B7Y0). MAAS also needs a dedicated XCC3 account with administrator rights and a password that does not expire, and its rack controllers must be able to reach the XCC3 management network. The Small design does not use MAAS and needs none of this.

Faster delivery with Lenovo Top Choice Express

Lenovo Top Choice Express is a configure-to-order program built on a curated set of the most popular components. Customers still choose the configuration they need, but the servers are built and fully integrated in Lenovo manufacturing with simpler configuration, instant pricing and significantly shorter lead times than a standard configure-to-order build.

The reference configurations in this brief are designed around Top Choice Express. The SR630 V4 and SR650 V4 chassis, processors, memory, NVMe backplanes, UEFI presets and the IPMI-over-LAN setting

all come from the Top Choice Express component list, which helps customers move from purchase order to a running cloud sooner. Top Choice Express status for the drives, boot kits, network adapters and power supplies will be confirmed as their part numbers are finalized.

Uses Cases

Migrating off proprietary virtualization

Organizations facing steep hypervisor renewals can consolidate virtualized workloads onto KVM-based Canonical OpenStack running on ThinkSystem SR630 V4 and SR650 V4 servers. A simple per-node subscription replaces per-core and per-VM licensing, so costs become predictable as the environment grows. Standard OpenStack APIs, Terraform, and Ansible give operations teams a familiar automation path, and the open stack removes the exit costs that made the previous platform hard to leave.

Sovereign cloud for the public sector and regulated industries

Government agencies, financial institutions, and service providers need self-service infrastructure-as-a-service while keeping full control of data location and operations. Canonical OpenStack delivers multi-tenant IaaS with quotas and metering on hardware the organization owns, in facilities it chooses, running transparent open-source code. FIPS certified cryptographic modules for the Ubuntu host, DISA-STIG and CIS hardening, and up to 15 years of security maintenance available with Ubuntu Pro support long-lived compliance programs.

AI and machine learning close to the data

Data science teams want GPU capacity without moving sensitive data to a public cloud. The same OpenStack cloud provides GPU-accelerated instances through PCI passthrough, alongside Kubernetes clusters for training and inference, so AI workloads run next to the data and share the platform's automation, observability, and security controls with conventional VMs.

Workloads

Cloud functions (control, compute, and storage) are assigned to servers as roles and combined into one of three reference designs. Services, APIs, and the operating model are identical across all three:

- **Entry (hyperconverged)**: minimum 3 nodes; survives a node failure. Suited to edge sites, remote offices, and starter clouds.
- **Production (converged)**: minimum 12 nodes (3 infrastructure plus 9 cloud); survives the loss of a full availability zone or rack. Suited to general-purpose enterprise clouds.
- **Scale-out (disaggregated)**: 3 infrastructure nodes plus dedicated control, compute, and storage nodes; survives the loss of an availability zone. Suited to large clouds and service providers.

Supported workload profiles include general-purpose VMs, Kubernetes on OpenStack, GPU-accelerated instances, and telco network functions using SR-IOV, CPU pinning and huge pages on dedicated compute nodes, with DPDK subject to design validation.

Lenovo has sized a ready-to-quote starting point for each design, called Small, Medium and Large. They are described in the next section.

Reference Configurations

Each reference design maps to a Lenovo configuration built from ThinkSystem SR630 V4 and SR650 V4 servers. All three follow Canonical's sizing guidelines, and all three run the same software, so customers can start small and grow without learning a new platform. The core of every server (chassis, processors, memory, NVMe backplanes and BIOS preset) uses Lenovo Top Choice Express components, which shortens delivery compared with a standard configure-to-order build. Table 1 compares them.

Table 1. Reference configurations at a glance

	Small	Medium	Large
Reference design	Entry (hyperconverged)	Production (converged)	Scale-out (disaggregated)
Nodes	3	12	36, plus 3 optional GPU nodes
Servers	3x SR630 V4	3x SR630 V4 infrastructure 9x SR650 V4 cloud	SR630 V4 infrastructure, control and compute SR650 V4 storage and GPU
Survives the loss of	One node	One rack (availability zone)	One rack (availability zone)
vCPUs for VMs, after a failure	304	2,304	7,936
RAM for VMs, after a failure	About 1.7 TB	About 11.2 TB	About 32 TB
Ordering	Top Choice Express	Top Choice Express	Top Choice Express
Usable Ceph capacity	About 35 TB	About 138 TB	About 553 TB
Data network per node	2x 25GbE	2x 100GbE	2x 100GbE (compute) 4x 100GbE (storage)
Typical use	Edge sites, remote offices, starter clouds	General-purpose enterprise cloud	Large enterprise, sovereign cloud, service providers, AI

The VM figures already hold back enough capacity to ride through the failure listed, and assume a 2:1 CPU overcommit with no memory overcommit. Usable storage allows for Ceph keeping three copies of every object, one per availability zone, and keeps 25% of the cluster free so Ceph can rebuild data after a drive failure.

Configuration choices common to every size

- **Boot**: two 960GB M.2 NVMe drives mirrored in RAID-1, kept separate from the data drives.
- **Storage**: enterprise read-intensive NVMe drives with power-loss protection, connected straight to the processors through the onboard NVMe ports. No RAID controller or HBA sits in the path. Each drive runs as one Ceph OSD, and every node in a storage tier uses the same number and size of drives. Unbalanced nodes or racks reduce usable capacity, because each availability zone holds one full copy of the data and the smallest one sets the limit.
- **Memory**: every memory channel on both processors is populated equally, so memory stays balanced. Nodes with one DIMM per channel run at 6400 MHz. Nodes that reach 2 TB use two 64GB DIMMs per channel, which runs at 5200 MHz and keeps the build on Top Choice Express parts.
- **Top Choice Express**: chassis, processors, memory, NVMe backplanes and UEFI presets come from the Top Choice Express component list for each platform. The remaining components will be confirmed as part numbers are finalized.
- **Network**: one OCP 3.0 adapter and one PCIe adapter in every node, bonded with LACP to a redundant switch pair, so the loss of an adapter or a switch does not take a node offline. Use passive copper (DAC) cables or optical transceivers at 100GbE in SR630 V4 nodes, because active optical cables tighten that server's thermal limits.
- **Management**: a dedicated 1GbE XClarity Controller 3 port on a separate management switch, and the optional XCC3 Premier upgrade for remote console and virtual media. In the Medium and Large designs, the MAAS rack controllers on the infrastructure nodes must reach this network; avoid multi-path routing on it, because it disrupts IPMI.

Small: Entry hyperconverged

Three 1U servers each run control, compute and storage. The Small design is installed manually: Ubuntu Server is installed on each server, and Canonical OpenStack is then deployed with its snap-based installer. It does not include MAAS, a separate observability stack or Landscape, and all cloud traffic shares a single network. These trade-offs keep the footprint to three servers and make Small the entry-level design; the Medium and Large designs add dedicated infrastructure nodes and separate networks. Two 32-core processors per node give Ceph enough CPU without crowding out tenant VMs, and the 10-bay NVMe backplane ships with six drives so capacity can grow in place. With three nodes and three data copies, a failed node leaves Ceph running on two copies until it is replaced, which suits edge and starter clouds well.

Table 2. Small configuration (3 nodes)

Component	Primary Role
Server	ThinkSystem SR630 V4 (1U)
Processors	2x Intel Xeon 6530P, 32 cores, 2.3 GHz
Memory	1 TB (16x 64GB TruDDR5 6400 MHz RDIMMs)
Boot	2x 960GB M.2 NVMe, RAID-1
Ceph storage	6x 7.68TB NVMe (10-bay NVMe backplane, 4 bays free)
Network	2x 25GbE on two adapters, plus 1GbE XCC3 management
Switching	2x 25GbE top-of-rack switches (MLAG), 1x 1GbE management switch

Medium: Production converged

Three small infrastructure nodes run MAAS, Juju, COS and Landscape. Nine 2U cloud nodes share control, compute and storage, three per rack, with each rack acting as its own availability zone. Each cloud node pairs two 64-core processors with 2 TB of memory, which leaves about 384 vCPUs and 1.9 TB of RAM per node for tenant VMs after Ceph and the control plane take their share. The SR650 V4 leaves eight drive bays and spare PCIe slots free in every cloud node, so customers can add storage or accelerators later. 100GbE networking keeps Ceph rebuilds fast on drive-dense nodes.

Table 3. Medium configuration (12 nodes)

Component	Infrastructure node (x3)	Cloud node (x9)
Server	ThinkSystem SR630 V4 (1U)	ThinkSystem SR650 V4 (2U)
Processors	2x Intel Xeon 6520P, 24 cores	2x Intel Xeon 6767P, 64 cores
Memory	512 GB (16x 32GB RDIMMs)	2 TB (32x 64GB RDIMMs)
Boot	2x 960GB M.2 NVMe, RAID-1	2x 960GB M.2 NVMe, RAID-1
Data drives	4x 3.84TB NVMe for management VMs, logs and metrics (10-bay backplane)	8x 7.68TB NVMe Ceph OSDs (16 bays, 8 free)
Network	2x 25GbE on two adapters	2x 100GbE on two adapters

Each of the three racks gets a 100GbE leaf switch pair and a 1GbE management switch, and two spine switches running BGP with ECMP tie the racks together. The control plane services sit behind virtual IPs spread across the racks, so the API, internal and provider networks must span all three racks. The fabric has to provide those networks, and Canonical defines the exact network layout during design.

Large: Scale-out disaggregated

Every role has its own servers, so compute and storage grow independently. Compute nodes pack two 64-core processors and 2 TB of memory into 1U. Storage nodes use all 24 front NVMe bays of the SR650 V4, with separate 100GbE networks for client access and replication. The design scales to roughly 200 to 250 nodes in a single cloud; beyond that, customers add regions.

Table 4. Large configuration (36 nodes)

Component	Infrastructure (x3)	Control (x3)	Compute (x24)	Storage (x6)
Server	SR630 V4	SR630 V4	SR630 V4	SR650 V4
Processors	2x Xeon 6520P 24 cores	2x Xeon 6530P 32 cores	2x Xeon 6767P 64 cores	2x Xeon 6767P 64 cores
Memory	512 GB	512 GB	2 TB (32x 64GB)	512 GB
Boot	2x 960GB M.2, RAID-1	2x 960GB M.2, RAID-1	2x 960GB M.2, RAID-1	2x 960GB M.2, RAID-1
Data drives	4x 3.84TB NVMe	2x 3.84TB NVMe for databases and queues	Optional 2x 3.84TB NVMe for local ephemeral disks	24x 15.36TB NVMe Ceph OSDs
Network	2x 25GbE	2x 25GbE	2x 100GbE	4x 100GbE (access and replication)

Large adds 11,904 vCPUs and about 48 TB of RAM across the compute tier, with about 2.2 PB of raw NVMe in the storage tier. Customers who want AI close to their data can add one SR650 V4 GPU node per availability zone, each with two 32-core Xeon 6730P processors, two double-wide data-center GPUs and 1 TB of memory, placed in its own host aggregate. Specific GPU models will be listed in a future revision.

Growing the cloud

Small grows by filling its empty drive bays and then by adding nodes in sets of three. Medium grows by adding cloud nodes three at a time, one per rack, and filling the free bays in each cloud node. Moving from Medium to Large means adding dedicated compute and storage nodes and moving control services onto their own servers, which is done as a planned migration with Canonical's automation rather than a new deployment. Moving from Small to Medium is a new deployment with workload migration, because Medium adds automated bare-metal provisioning, dedicated infrastructure nodes, separate networks and rack-level availability zones; the SR630 V4 servers can be reused in the new cloud.

Business Outcome

Canonical OpenStack on Lenovo ThinkSystem SR630 V4 and SR650 V4 delivers value in several measurable areas:

- **Lower, predictable TCO**: no per-VM or per-core licensing; one per-node subscription covers the software stack, with Ceph capacity beyond the amount included in Ubuntu Pro priced per terabyte.
- **Freedom of choice**: upstream OpenStack APIs, standard KVM, OVN, and Ceph, and industry-standard Lenovo ThinkSystem SR630 V4 and SR650 V4 hardware remove lock-in.
- **Operational simplicity**: model-driven automation handles deployment, scaling, upgrades, and observability, reducing the specialist effort a cloud demands.
- **Faster time to value**: reference configurations built around Lenovo Top Choice Express core components ship with shorter lead times than standard configure-to-order builds.
- **Enterprise resilience**: a highly available design tolerates the loss of a node, a rack, or an availability zone.
- **Long-term security**: up to 15 years of security maintenance per Ubuntu LTS release (10 years with Ubuntu Pro and 5 more with the Legacy add-on), plus compliance tooling.

Specific TCO and performance figures for this solution have not yet been published. Contact your Lenovo representative for a sizing and cost comparison based on your environment.

Bill of Materials

The following table lists the bill of materials for the Medium configuration, the design most enterprise clouds start from. Quantities are totals for the whole cloud. Processors, memory, chassis, backplanes, the UEFI preset and the IPMI-over-LAN option are all Top Choice Express components, and DCSC selects the heatsinks automatically; the TCE column marks the items whose Top Choice Express status will be confirmed when their part numbers are finalized in a future revision. Contact your Lenovo representative for the Small and Large bills of materials.

Table 5. Bill of Materials, Medium Configuration

Part Number	Description	Quantity	TCE
Infrastructure Nodes (3)			
7DG9CTO1WW	ThinkSystem SR630 V4, 3-year warranty (base chassis feature C1XE, 10x 2.5-inch)	3	Yes
4XG7B17661	ThinkSystem SR630 V4 Intel Xeon 6520P 24C 210W 2.4GHz Processor w/o Fan v2	6	Yes
4X77A90964	ThinkSystem 32GB TruDDR5 6400MHz (1Rx4) RDIMM	48	Yes
C21X	ThinkSystem 1U V4 10x2.5" NVMe Gen5 Backplane	3	Yes
C1XJ	ThinkSystem 1U V4 Performance Heatsink (derived by DCSC)	6	Yes
C9U9	ThinkSystem Virtualization, Max Performance UEFI preset	3	Yes
B7Y0	Enable IPMI-over-LAN (required by MAAS)	3	Yes
TBD	3.84TB read-intensive NVMe 2.5-inch SSD	12	To confirm
TBD	M.2 RAID-1 boot kit with 2x 960GB M.2 NVMe drives	3	To confirm
TBD	Dual-port 25GbE SFP28 OCP 3.0 adapter	3	To confirm
TBD	Dual-port 25GbE SFP28 PCIe adapter	3	To confirm
TBD	Titanium hot-swap power supply, sized in DCSC	6	To confirm
TBD	XClarity Controller 3 Premier upgrade (optional)	3	To confirm
Cloud nodes (9)			
7DGDCTO1WW	ThinkSystem SR650 V4, 3-year warranty (base chassis feature C3QK, 24x 2.5-inch)	9	Yes
4XG7B03692	ThinkSystem SR650/a V4 Intel Xeon 6767P 64C 350W 2.4GHz Processor w/o Fan	18	Yes
4X77A90966	ThinkSystem 64GB TruDDR5 6400MHz (2Rx4) RDIMM	288	Yes
C46P	ThinkSystem 2U V4 8x2.5" NVMe Backplane	18	Yes
Derived	Processor heatsink, selected automatically by DCSC per SR650 V4 thermal rules	18	Yes
C9U9	ThinkSystem Virtualization, Max Performance UEFI preset	9	Yes
B7Y0	Enable IPMI-over-LAN (required by MAAS)	9	Yes
TBD	7.68TB read-intensive NVMe PCIe 5.0 2.5-inch SSD	72	To confirm
TBD	M.2 RAID-1 boot kit with 2x 960GB M.2 NVMe drives	9	To confirm
TBD	Dual-port 100GbE QSFP56 OCP 3.0 adapter	9	To confirm
TBD	Dual-port 100GbE QSFP56 PCIe adapter	9	To confirm
TBD	Titanium hot-swap power supply, sized in DCSC	18	To confirm
TBD	XClarity Controller 3 Premier upgrade (optional)	9	To confirm
Software			

Part Number	Description	Quantity	TCE
TBD	Ubuntu Pro + Support (Infra, 24/7), per physical node	12	Not applicable

Top Choice Express ship dates are estimates, and availability of individual components can vary by region. Check the Top Choice Express column in the Lenovo Data Center Solution Configurator (DCSC) for your region when building the quote. Ceph capacity beyond the amount included in Ubuntu Pro is priced per terabyte by Canonical.

Conclusion

Canonical OpenStack on Lenovo ThinkSystem SR630 V4 and SR650 V4 delivers an open, secure, and cost-predictable private cloud that combines upstream OpenStack, model-driven automation, and long-term security maintenance with Lenovo's enterprise-class infrastructure. The architecture scales from edge to data center, and customers can operate it themselves with Canonical support or have Canonical manage it for them. For organizations moving away from costly proprietary virtualization, or building sovereign and AI-ready clouds, it offers a practical, low-risk path to cloud agility on their own terms.

With Small, Medium and Large reference configurations built on Lenovo Top Choice Express, customers can size, quote and receive their cloud hardware faster.

For More Information

To learn more about Canonical OpenStack on Lenovo ThinkSystem SR630 V4 and SR650 V4, contact your Lenovo representative or Lenovo Business Partner, or visit the resources below. References:

- Canonical OpenStack:
  https://ubuntu.com/openstack
- Canonical OpenStack documentation:
  https://canonical.com/openstack/docs/stable-2024.1/
- Ubuntu Pro:
  https://ubuntu.com/pro
- Canonical Charmed Kubernetes on Lenovo ThinkSystem for AI Development:
  https://lenovopress.lenovo.com/lp1415-canonical-charmed-kubernetes-on-thinksystem-for-ai-development
- Lenovo OS Interoperability Guide:
  https://lenovopress.lenovo.com/osig
- MAAS (Metal as a Service):
  https://canonical.com/maas

Authors

Jessie Lacome is a Platform Solutions Product Manager at Lenovo with extensive experience in enterprise virtualization and modern infrastructure platforms. He brings a strong consulting and solutions architecture background, supporting customers across on-premises and hybrid environments. Jessie has previously held roles at Salesforce, Citrix, and Dell, where he focused on virtualization and infrastructure solutions.

Related product families

Product families related to this document are the following:

- ThinkSystem SR630 V4 Server
- ThinkSystem SR650 V4 Server

Notices

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

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

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

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

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

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

© Copyright Lenovo 2026. All rights reserved.

This document, LP2568, was created or updated on November 17, 2026.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
  https://lenovopress.lenovo.com/LP2568
- Send your comments in an e-mail to:
  comments@lenovopress.com

This document is available online at https://lenovopress.lenovo.com/LP2568.

Trademarks

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

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

The following terms are trademarks of other companies:

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

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

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