

The Lenovo logo, consisting of the word "Lenovo" in white lowercase letters on a black rectangular background.

# Lenovo Database Reference Architecture for IBM Db2 on IBM Cloud Private

Last update: **8 February 2019**  
Version 1.0

---

**Describes the reference architecture for high performance infrastructure for IBM Cloud Private**

---

**Solution based on the powerful, versatile Lenovo ThinkSystem SR650 servers powered by Intel Xeon Scalable Processors**

---

**Deployment considerations for high-performance, cost-effective and scalable solutions**

---

**Uses Intel NVMe storage and Lenovo network devices to deliver very high performance**

Lenovo  
IBM  
Intel



# Table of Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Business problem and business value</b>	<b>2</b>
2.1	Business problem	2
2.2	Business value	2
<b>3</b>	<b>Requirements</b>	<b>3</b>
3.1	Functional requirements	3
3.2	Non-functional requirements	3
<b>4</b>	<b>Architectural overview</b>	<b>4</b>
<b>5</b>	<b>Component model</b>	<b>5</b>
5.1	IBM Cloud Private Components	5
<b>6</b>	<b>Operational model</b>	<b>7</b>
6.1	Hardware description	7
6.1.1	Lenovo ThinkSystem SR650 Server	7
6.1.2	Lenovo RackSwitch G8052	8
6.1.3	Lenovo RackSwitch G8272	8
6.1.4	Lenovo RackSwitch NE2572	9
6.1.5	Lenovo RackSwitch NE10032	10
6.2	Cluster nodes	10
6.2.1	Node Types	11
6.2.2	Node Hardware Configuration	12
6.2.3	Cluster Sizes	14
6.3	IBM Cloud Private software stack	14
6.4	Systems management	15
6.5	Networking	16
6.5.1	Data Network	17
6.5.2	Hardware Management network	17
6.5.3	10Gb and 25Gb Data Network Configurations	18
<b>7</b>	<b>Deploying VMware vSphere</b>	<b>19</b>

7.1	vSphere Configuration .....	19
7.2	vSphere Distributed Switch .....	21
7.2.1	Network Architecture .....	21
7.2.2	Management Network .....	22
7.3	VMware vSAN configuration .....	23
<b>8</b>	<b>Deploying IBM Cloud Private .....</b>	<b>24</b>
8.1	Cluster Monitoring .....	24
8.1.1	Logging .....	24
8.1.2	Monitoring and Alerting .....	25
8.1.3	Metering .....	26
8.2	Network Configurations - IBM Cloud Private .....	27
<b>9</b>	<b>Deploying IBM Db2 .....</b>	<b>29</b>
9.1	Estimating cluster size .....	<b>Error! Bookmark not defined.</b>
9.2	High availability considerations .....	29
<b>10</b>	<b>Acknowledgements .....</b>	<b>31</b>
	<b>Resources .....</b>	<b>32</b>
	<b>Document history .....</b>	<b>33</b>

# 1 Introduction

---

This document describes the reference architecture for the IBM Cloud Private solution. It provides a predefined and optimized hardware infrastructure for high performance implementation of the IBM Cloud Private software. It also provides an example of deploying the IBM Db2 package on IBM Cloud Private. This reference architecture provides planning, design considerations, and best practices for implementing IBM Cloud Private with Lenovo and Intel products.

Lenovo, IBM and Intel teams worked together on this document and the reference architecture described herein was developed and validated in a joint engineering project.

Container-based architectures have become increasingly attractive to modern businesses that want to digitally transform their data centers with a modern application and technology infrastructure. Containers provide a lightweight complement to virtual machines (VMs) that are portable, run anywhere, are fast and easy to deploy, and can support both traditional apps and microservices. Containers are also attractive from a business perspective because they can be used with VMs to reduce hardware, licensing, and maintenance costs, accelerate time to market, increase productivity, and enable hybrid clouds.

The Lenovo validated design for IBM Cloud Private enables IT organizations to deploy a container solution cost effectively, quickly, and without adding complexity. The combination of IBM Cloud Private and VMware vSAN on Lenovo ThinkSystem hardware can provide an economical, enterprise-grade way forward for digital transformation. In addition, the validated design's performance-optimized infrastructure is built on Intel Xeon Scalable processors, Intel 3D NAND SSDs, Intel Optane DC SSDs, and the Intel Ethernet 700 Series network adapters.

The intended audience of this document is IT professionals, technical architects, sales engineers, and consultants to assist in planning, designing, and implementing the big data solution with Lenovo hardware. It is assumed that you are familiar with virtualization and data warehouse concepts. For more information, see "Resources" on page 32.

## 2 Business problem and business value

---

This section describes business challenges faced by big data environments and the value provided by the IBM Cloud Private solution used to address the business challenges.

### 2.1 Business problem

Cloud computing has become pervasive for providing data analytic and storage capabilities in a managed environment that can be scaled elastically as business needs change. The public cloud, available as a service, provides efficiency and rapid provisioning for business applications; but organizations in many cases cannot run their applications on the public cloud. Most of the time they are held back by regulatory and legal requirements. Others do not want their strategic assets to leave their organization. Private cloud can bring the benefits of a public cloud to your on-premises infrastructure and keep it behind your firewall. A private cloud is like a fenced-in backyard with a gate to surround properties and public spaces. It can give you many of the benefits of a public cloud with the additional control and security of dedicated resources.

### 2.2 Business value

IBM Cloud Private offers similar benefits to a public cloud but is dedicated to a single organization, allowing customization of the environment to meet personalized needs and security requirements. It can be managed internally or by a third party and it can be hosted behind the company firewall or externally. This platform offers public cloud advantages while providing the increased performance, predictable costs, tighter security and flexible management options found in a private cloud.

The IBM Cloud Private offering extends directly from IBM's experience with cloud-native infrastructure and applications. The principles that guide the IBM Cloud Private platform are:

- Enterprise-focused platform using microservices - many of the advantages of microservices come from resource isolation, scale up and scale down, and lightweight, portable movement of application workloads.
- Application services - runtimes, middleware, data and other value-added services can be added to cloud-native applications or connected to existing applications. IBM Cloud Private makes it easier to stand up an elastic runtime that is based on a variety of workloads.

# 3 Requirements

---

The functional and non-functional requirements for this reference architecture are described in this section.

## 3.1 Functional requirements

A private cloud solution for big data analytics supports the following key functional requirements:

- Ability to optimize applications for the cloud
- Open data centers to work with cloud services
- Create new cloud native applications

## 3.2 Non-functional requirements

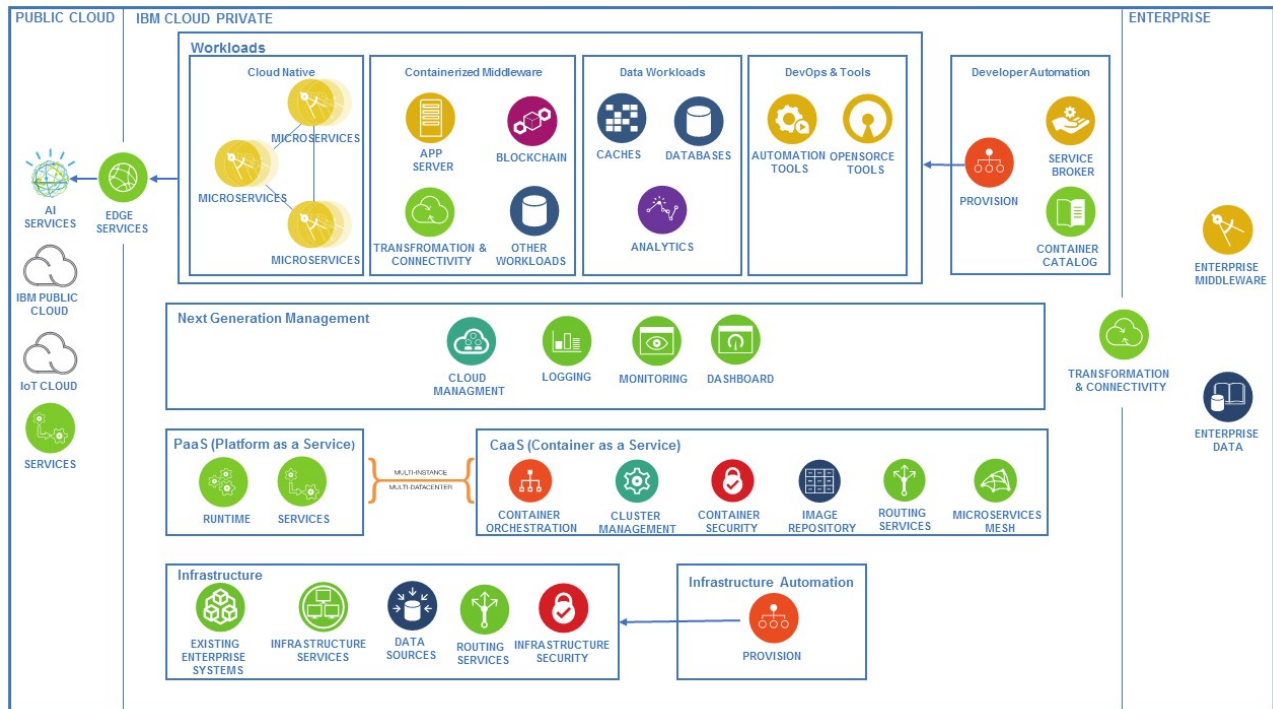
Customers require their cloud private infrastructure to be easy, dependable, and fast. The following non-functional requirements are key:

- Easy:
  - Ease of development
  - Easy management at scale
  - Advanced job management
  - Easy to access data by various user types
- Dependable:
  - Data protection with snapshot and mirroring
  - Automated self-healing
  - Insight into software/hardware health and issues
  - High availability (HA) and business continuity
- Fast:
  - Superior performance
  - Scalability
- Secure and governed:
  - Strong authentication and authorization
  - Physical and firewall security
  - Data confidentiality and integrity

# 4 Architectural overview

The IBM private cloud solution is based on a flexible and scalable architecture. The primary hardware building block is the worker node implemented on Lenovo ThinkSystem SR650 servers. A cluster of SR650 servers are connected together to meet the desired total compute and storage capacity required to deliver the best performance for IBM Cloud Private solutions.

Figure 1 shows the architecture overview of the IBM Cloud Private reference architecture that uses Lenovo ThinkSystem hardware infrastructure.



**Figure 1. IBM Cloud Private architecture overview.**

The private cloud architecture delivers several key functions:

- Container orchestration based on Kubernetes - this layer provides cluster management, security capabilities, image repositories, routing services, and microservices mesh.
- A PaaS layer to provide higher-level runtimes and service bindings for an easier development experience.
- The CaaS and PaaS layer, which sit over an infrastructure layer to provide compute through virtual machines, network, storage, and security.
- Automation and orchestration for the underlying infrastructure to provide predefined infrastructure templates to create repeatable patterns.
- Monitoring for container-based applications to provide logging, dashboards, and automation. It supports network and storage policy-based controls for application isolation and security, and automated application health checking and recovery from failures.
- The ability to run containerized workloads for several patterns, such as cloud-native, data workloads, integration workloads, tool workloads, and some middleware, such as Java Application Server.
- A workloads catalog that you can provision via containers to automate the development experience.

# 5 Component model

This section describes the high-level component model of the IBM Cloud Private solution shown in the following diagram. By supporting Kubernetes and Cloud Foundry, IBM Cloud Private provides choices in application development such as content that can be containerized, tools for end-to-end automation, and management tools.

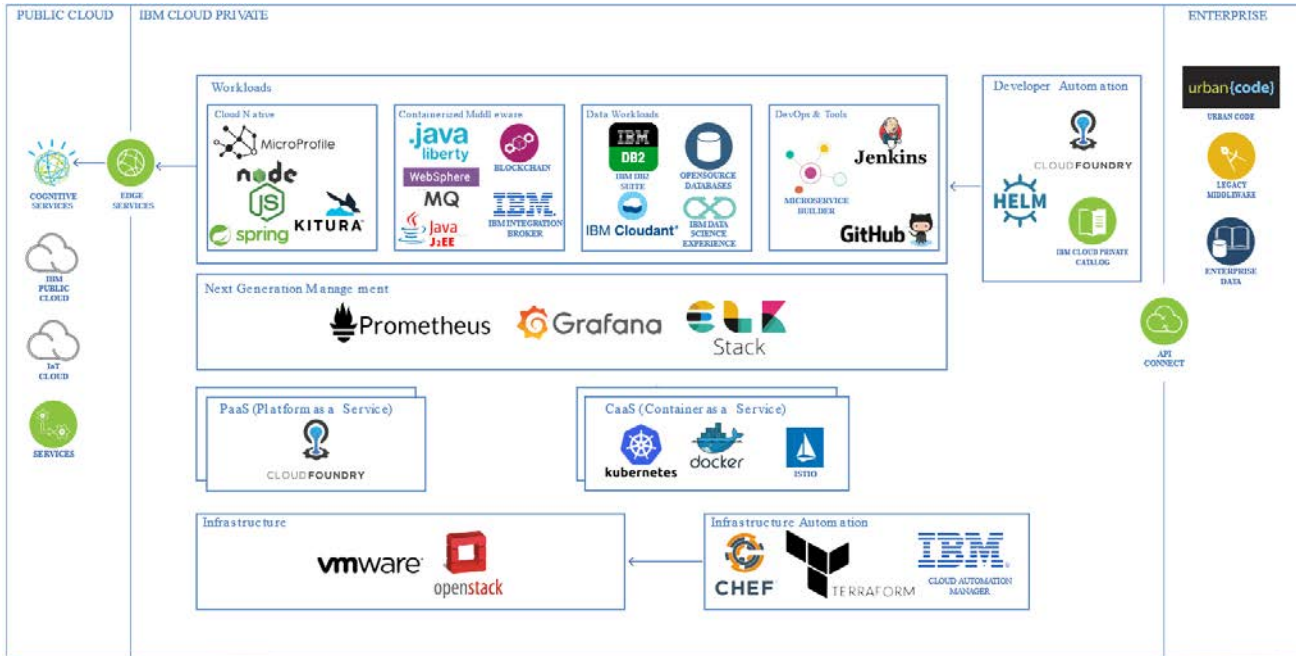


Figure 2. IBM Cloud Private component model overview.

## 5.1 IBM Cloud Private Components

Several primary components of the IBM Cloud Private solution are described below:

- **Kubernetes** - IBM Cloud Private provides an open container platform based on Kubernetes for automating deployment, scaling, and management of containerized applications. It groups containers that make up an application into logical units for easy management and discovery.
- **Docker** - Docker containers provide lightweight, standalone, executable software packages with everything needed to run an application, regardless of the infrastructure or operating system. In IBM Cloud Private, a private Docker registry integrates with the Docker registry V2 API to provide a local registry service that functions in the same way as the cloud-based registry service, Docker Hub. This local registry has all the same features as Docker Hub, but you can restrict which users can view or pull images from this registry.
- **Helm** - Helm is the Kubernetes native package management system and is used for application management inside an IBM Cloud Private cluster. The Helm GitHub community curates and continuously expands a set of tested and preconfigured Kubernetes applications. You can add items



from this catalog of stable applications to your cluster from the management console. Installing this Helm community catalog provides an extra 80+ Kubernetes applications that are ready for deployment in your cluster.

- **IBM Cloud Automation Manager** - IBM Cloud Automation Manager is a multi-cloud, self-service management platform running on IBM Cloud Private that empowers developers and administrators to meet business demands. This platform allows you to efficiently manage and deliver services through end-to-end automation while enabling developers to build applications aligned with enterprise policies.
- **IBM Cloud Private Catalog** - IBM Cloud Private provides an easy to use, extend, and compose Catalog of IBM and third-party content. The Catalog provides a centralized location from which you can browse for and install packages in your cluster. Packages for additional IBM products are available from curated repositories that are included in the default IBM Cloud Private repository list which is available via an internet connection.
- **IBM Db2** - Db2 Database and Db2 Warehouse includes products for operational databases, data warehouses, data lakes and fast data. These databases and data warehouses support best-in-class and open-source software.
- **Cloud Foundry** - Cloud Foundry is open source software that makes it faster and easier to build, test, deploy and scale applications, providing a choice of clouds, developer frameworks, and application services. It is available through a variety of private cloud distributions and public cloud instances.
- **VMware** - is a leading virtual machine software provider which powers a large number of worldwide on-premises and cloud platforms. VMware virtualization increases utilization of hardware resources while giving easy provisioning of software platforms onto the hardware.
- **Grafana** - is open source software that allows you to query, visualize, create alerts on data metrics where ever they are stored. One can create, explore, and share dashboards that provide an easy to use analytics platform for data metrics.

# 6 Operational model

---

This section describes the hardware infrastructure aspects of the IBM Cloud Private reference architecture. To support different customer environments, different configurations are provided for supporting different amounts of data sizes and performance levels.

## 6.1 Hardware description

This reference architecture uses Lenovo servers SR650 (2U) servers and Lenovo RackSwitch G8052 and G8272 top of rack switches.

### 6.1.1 Lenovo ThinkSystem SR650 Server

The Lenovo ThinkSystem SR650 is an ideal 2-socket 2U rack server for small businesses up to large enterprises that need industry-leading reliability, management, and security, as well as maximizing performance and flexibility for future growth. The SR650 server is particularly suited for big data applications due to its rich internal data storage, large internal memory and selection of high performance Intel processors. It is also designed to handle general workloads, such as databases, virtualization and cloud computing, virtual desktop infrastructure (VDI), enterprise applications, collaboration/email, and business analytics.

The SR650 server supports:

- Up to two Intel Xeon Scalable Processors
- Up to 1.5 TB 2666 MHz TruDDR4 memory (support for up to 3 TB is planned for future),
- Up to 24x 2.5-inch or 14x 3.5-inch drive bays with an extensive choice of NVMe PCIe SSDs, SAS/SATA SSDs, and SAS/SATA HDDs
- Flexible I/O Network expansion options with the LOM slot, the dedicated storage controller slot, and up to 6x PCIe slots



**Figure 3.** Lenovo ThinkSystem SR650

Combined with the Intel Xeon Scalable Processors (Bronze, Silver, Gold, and Platinum), the Lenovo SR650 server offers an even higher density of workloads and performance that lowers the total cost of ownership (TCO). Its pay-as-you-grow flexible design and great expansion capabilities solidify dependability for any kind of workload with minimal downtime.

The SR650 server provides high internal storage density in a 2U form factor with its impressive array of workload-optimized storage configurations. It also offers easy management and saves floor space and power consumption for most demanding use cases by consolidating storage and server into one system.

This reference architecture recommends the storage-rich ThinkSystem SR650 for the following reasons:

- **Storage capacity:** Each of the 14 configured 3.5-inch drives has raw capacity up to 10 TB and each, providing for 140 TB of raw storage per node and over 2000 TB per rack.
- **Performance:** This hardware supports the latest Intel Xeon Scalable processors and TruDDR4 Memory.
- **Flexibility:** Server hardware uses embedded storage, which results in simple scalability (by adding nodes).
- **PCIe slots:** Up to 7 PCIe slots are available if rear disks are not used, and up to 3 PCIe slots if the Rear HDD kit is used. They can be used for network adapter redundancy and increased network throughput.
- **Higher power efficiency:** Titanium and Platinum redundant power supplies can deliver 96% (Titanium) or 94% (Platinum) efficiency at 50% load.
- **Reliability:** Outstanding reliability, availability, and serviceability (RAS) improve the business environment and helps save operational costs

For more information, see the Lenovo ThinkSystem SR650 Product Guide:

<https://lenovopress.com/lp0644-lenovo-thinksystem-sr650-server>

### 6.1.2 Lenovo RackSwitch G8052

The Lenovo networking RackSwitch G8052 (as shown in Figure 4) is an Ethernet switch that is designed for the data center and provides a simple network solution. The Lenovo RackSwitch G8052 offers up to 48x 1 GbE ports and up to 4x 10 GbE ports in a 1U footprint. The G8052 switch is always available for business-critical traffic by using redundant power supplies, fans, and numerous high-availability features.



**Figure 4. Lenovo RackSwitch G8052**

Lenovo RackSwitch G8052 has the following characteristics:

- A total of **48x 1 GbE** RJ45 ports
- **Four 10 GbE** SFP+ ports
- Low 130W power rating and variable speed fans to reduce power consumption

For more information, see the Lenovo RackSwitch G8052 Product Guide:

<https://lenovopress.com/tips1270-lenovo-rackswitch-g8052>

### 6.1.3 Lenovo RackSwitch G8272

Designed with top performance in mind, Lenovo RackSwitch G8272 is ideal for today's big data, cloud, and optimized workloads. The G8272 switch offers up to 72 10 Gb SFP+ ports in a 1U form factor and is

expandable with four 40 Gb QSFP+ ports. It is an enterprise-class and full-featured data center switch that deliver line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data center grade buffers keep traffic moving. Redundant power and fans and numerous HA features equip the switches for business-sensitive traffic.

The G8272 switch (as shown in Figure 5) is ideal for latency-sensitive applications, such as client virtualization. It supports Lenovo Virtual Fabric to help clients reduce the number of I/O adapters to a single dual-port 10 Gb adapter, which helps reduce cost and complexity. The G8272 switch supports the newest protocols, including Data Center Bridging/Converged Enhanced Ethernet (DCB/CEE) for support of FCoE and iSCSI and NAS.



**Figure 5: Lenovo RackSwitch G8272**

The enterprise-level Lenovo RackSwitch G8272 has the following characteristics:

- 48 x SFP+ 10GbE ports plus 6 x QSFP+ 40GbE ports
- Support up to 72 x 10Gb connections using break-out cables
- 1.44 Tbps non-blocking throughput with very low latency (~ 600 ns)
- Up to 72 1Gb/10Gb SFP+ ports
- OpenFlow enabled allows for easily created user-controlled virtual networks

For more information, see the RackSwitch G8272 Product Guide: [lenovopress.com/tips1267](http://lenovopress.com/tips1267).

### 6.1.4 Lenovo RackSwitch NE2572

The Lenovo ThinkSystem NE2572 RackSwitch is designed for the data center and provides 10 Gb/25 Gb Ethernet connectivity with 40 Gb/100 Gb Ethernet upstream links. It is ideal for big data workload solutions and is an enterprise class Layer 2 and Layer 3 full featured switch that delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data center-grade buffers help keep traffic moving, while the hot-swap redundant power supplies and fans (along with numerous high-availability software features) help provide high availability for business sensitive traffic.



### **Figure 6. Lenovo RackSwitch NE2572**

The NE2572 has the following characteristics:

- 48x SFP28/SFP+ ports that support 10 GbE SFP+ and 25 GbE SFP28 with AOC and DAC cabling
- 6x QSFP28/QSFP+ ports that support 40 GbE QSFP+ and 100 GbE QSFP28 optical transceivers with iAOC and DAC cabling
- QSFP28/QSFP+ ports can also be split out into two 50 GbE (for 100 GbE QSFP28), or four 10 GbE (for 40 GbE QSFP+) or 25 GbE (for 100 GbE QSFP28) connections by using breakout cables.

For more information, see the Lenovo RackSwitch NE2572 Product Guide:

<https://lenovopress.com/lp0608-lenovo-thinksystem-ne2572-rackswitch>

### **6.1.5 Lenovo RackSwitch NE10032**

The Lenovo ThinkSystem NE10032 RackSwitch that uses 100 Gb QSFP28 and 40 Gb QSFP+ Ethernet technology is specifically designed for the data center. It is ideal for today's big data workload solutions and is an enterprise class Layer 2 and Layer 3 full featured switch that delivers line-rate, high-bandwidth switching, filtering and traffic queuing without delaying data. Large data center-grade buffers help keep traffic moving, while the hot-swap redundant power supplies and fans (along with numerous high-availability features) help provide high availability for business sensitive traffic.

The NE10032 RackSwitch has 32x QSFP+/QSFP28 ports that support 40 GbE and 100 GbE optical transceivers, active optical cables (AOCs), and direct attach copper (DAC) cables. It is an ideal cross-rack aggregation switch for use in a multi rack cluster.



**Figure 7: Lenovo ThinkSystem NE10032 cross-rack switch**

For further information on the NE10032 switch, visit this link:

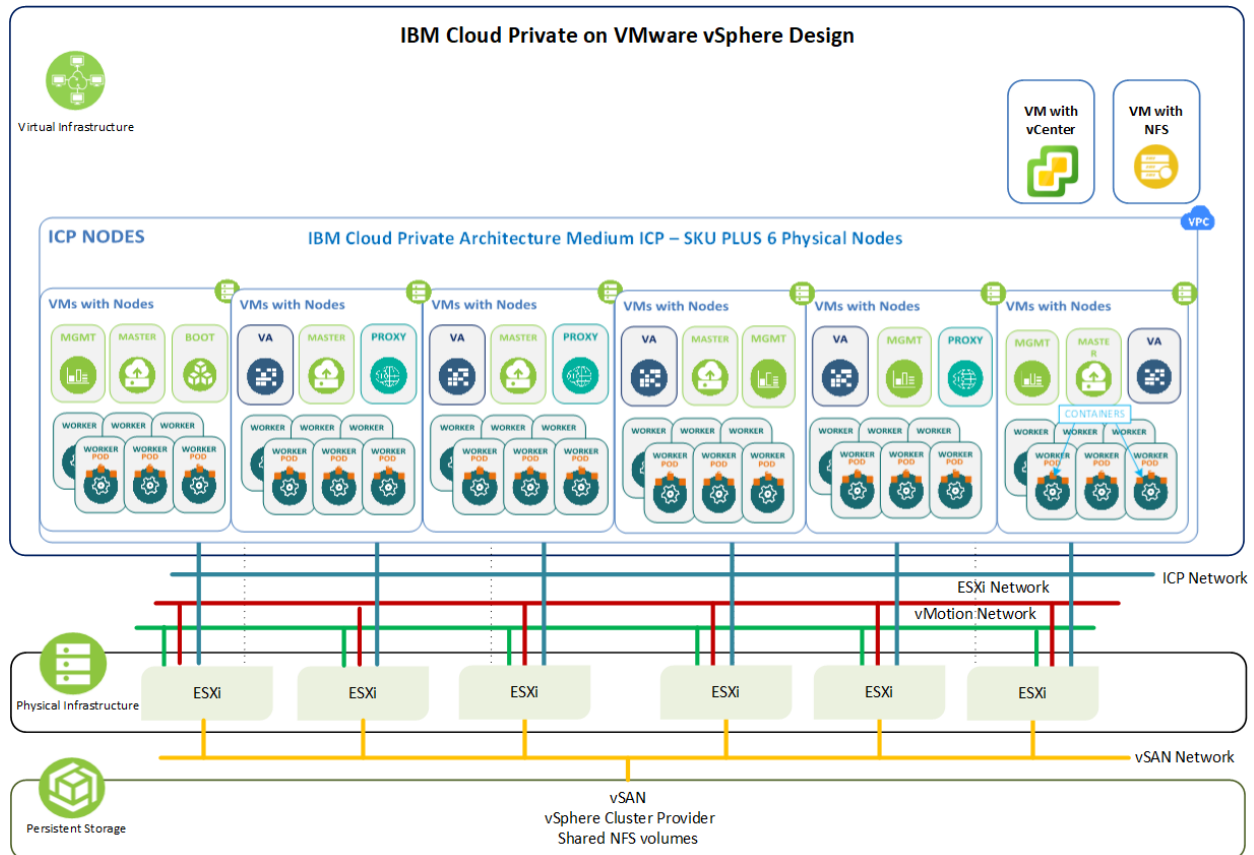
<https://lenovopress.com/lp0609-lenovo-thinksystem-ne10032-rackswitch>

## **6.2 Cluster nodes**

An IBM Cloud Private cluster is composed of certain node types and can scale from a small proof of concept (POC) cluster to large multi rack configurations. Nodes run various software services that support the full cloud private solution and provide high performance and high availability to the operating cluster. The physical nodes can be dedicated to each node type, or multiple Virtual Machines (VMs) can be configured on a node to run multiple node types. This reference architecture uses VMware Virtual Machines and vSAN ReadyNodes to demonstrate a hyper-converged infrastructure for a transformed, software-defined data center.

## 6.2.1 Node Types

Figure 11 below shows an IBM Cloud Private cluster with various node types used in this reference architecture. It includes 6 physical nodes, each containing Virtual Machines (VMs) created with VMware ESXi hypervisor which are used to implement the required node types.



**Figure 8. IBM Cloud Private cluster with medium-sized node configuration**

An IBM Cloud Private cluster has four main classes of nodes: boot, master, worker, and proxy, with two optional nodes—management and vulnerability advisor.

- **Boot Node**

A boot, or bootstrap, node is used for running installation, configuration, node scaling, and cluster updates. Only one boot node is required for any cluster. You can use a single node for both master and boot.

- **Master Node**

A master node provides management services and controls the worker nodes in a cluster. Master nodes host processes that are responsible for resource allocation, state maintenance, scheduling, and monitoring. Multiple master nodes can be configured in a high availability (HA) environment to allow for failover if the leading master host fails. Hosts that can act as the master are called master candidates.

- **Worker Node**

A worker node is a node that provides a containerized environment for running tasks. As demands

increase, more worker nodes can easily be added to your cluster to improve performance and efficiency (see the IBM Cloud Private cluster sizing guidelines in Table 4). A cluster can contain any number of worker nodes, but a minimum of one worker node is required.

- **Proxy Node**

A proxy node is a node that transmits external requests to the services created inside your cluster. Multiple proxy nodes can also be deployed in an HA environment to allow for failover if the leading proxy host fails. While you can use a single node as both master and proxy, it is best to use dedicated proxy nodes to reduce the load on the master node. A cluster must contain at least one proxy node if load balancing is required inside the cluster.

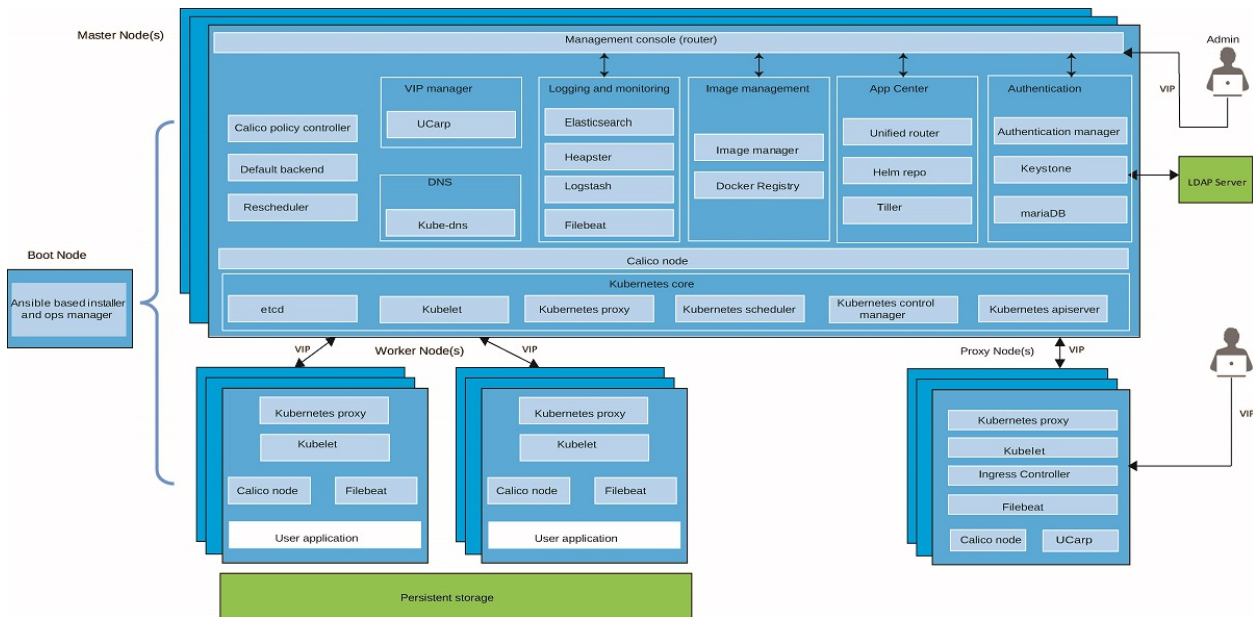
- **Management Node**

A management node is an optional node that only hosts management services such as monitoring, metering, and logging. By configuring dedicated management nodes, you can prevent the master node from becoming overloaded. You can enable a management node only during IBM Cloud Private installation.

- **Vulnerability Advisor Node**

A vulnerability advisor (VA) node is an optional node that is used for running the VA services, which can be resource-intensive. If you use the VA service, specify a dedicated VA node.

The following figure shows the interaction of the IBM Cloud Private nodes with their particular functionality.



**Figure 9. IBM Cloud Private node type function and interaction**

## 6.2.2 Node Hardware Configuration

Three node configurations for CPU core count, system memory, and storage are recommended for certain

types of cluster usage. Configurations of increasing performance are defined for a non-production development or test cluster (DEV/Test), a balanced configuration for best cost vs. performance (Base), and a higher performance configuration for increased node throughput (Plus). The Base configuration is a good starting point for most cluster usage in a production environment.

**Table 1. Dev/Test node configurations**

Component	Compute node configuration
Server	ThinkSystem SR650
Processor	2x Intel Xeon processors: Silver 4114, 2.20 GHz, 10 cores
Memory - base	12x 16GB 2,666MHz DDR4 (192GB)
Disk (OS)	Dual M.2 480GB SSD with RAID1
Disk (shared storage; worker nodes)	1x 375GB U.2 Intel P4800X NVMe SSD (2.5-in.) 3x 2TB U.2 Intel P4510 NVMe SSD (2.5-in.)
Hardware management network adapter	Integrated 1G BaseT XCC management controller, 1Gb LOM. dedicated or shared LAN port
Data network adapter	1x dual port X710-DA2 PCIe Adapter, 10Gb Ethernet

**Table 2. Base node Configuration**

Component	Compute node configuration
Server	ThinkSystem SR650
Processor	2x Intel Xeon processors: Gold 5120, 2.20GHz, 14 cores
Memory - base	12x 32GB 2,666MHz DDR4 (384GB)
Disk (OS)	Dual M.2 480GB SSD with RAID1
Disk (shared storage; worker nodes)	2x 375GB U.2 Intel P4800X NVMe SSD (2.5-in.) 6x 2TB U.2 Intel P4510 NVMe SSD (2.5-in.)
NVMe Switch Adapter	1x ThinkSystem 1610-4P NVMe Switch Adapter
Hardware management network adapter	Integrated 1G BaseT XCC management controller, 1Gb LOM. dedicated or shared LAN port
Data network adapter	1x dual port XXV710-DA2 PCIe Adapter, 25Gb Ethernet

**Table 3. Plus node configuration**

Component	Compute node configuration
Server	ThinkSystem SR650
Processor	2x Intel Xeon processors: Gold 6148, 2.50 GHz, 20 cores
Memory - base	24x 32GB 2,666MHz DDR4 (768GB)
Disk (OS)	Dual M.2 480GB SSD with RAID1
Disk (shared storage; worker nodes)	2x 375GB U.2 Intel P4800X NVMe SSD (2.5-in.) 8x 2TB U.2 Intel P4510 NVMe SSD (2.5-in.)
NVMe Switch Adapter	2x ThinkSystem 810-4P NVMe Switch Adapter 2x ThinkSystem 1610-4P NVMe Switch Adapter
Hardware management network adapter	Integrated 1G BaseT XCC management controller, 1Gb LOM. Dedicated or shared LAN port
Data network adapter	1x dual port XXV710-DA2 PCIe Adapter, 25Gb Ethernet



Other Processors with different core count and frequency are available for matching the compute intensity required by various workloads. The memory capacity can also be adjusted based on cost and performance considerations. Each node in the reference architecture has internal NVMe Solid State Drives.

The number of data nodes that are required within an IBM Cloud Private cluster depends on the client requirements. Such requirements might include the type of analytics workloads, the size of the user data, and the number of analytics tools being used.

### 6.2.3 Cluster Sizes

Recommended cluster sizes are shown in the following chart based on three sizes of physical node configurations (Dev/Test, Base, Plus).

**Table 4. Cluster Physical Node Sizes**

IBM Cloud Private Cluster Sizes	Small Cluster	Medium Cluster		Large Cluster	
	Max. 20 Worker VMs	Min. 5 Worker VMs	Max 70 Worker VMs	Min. 7 Worker VMs	Max. 150 Worker VMs
Dev/Test node configuration	4	5	17	-	-
Base node configuration	4	4	9	4	17
Plus node configuration	4	4	6	4	12

## 6.3 IBM Cloud Private software stack

The software stack deployed in this reference architecture is shown below.

**Table 5. Solution software stack**

Software Name	Version
VMware vSphere ESXi	ESXi 6.7 U1
VMware vCenter	VMware vCenter Server Appliance 6.7 U1
VMware vSAN	ESXi 6.7 U1
IBM Cloud Private Enterprise Edition	3.1.1
IBM Cloud Private Master Node	Red Hat Enterprise Linux (RHEL) v7.5
IBM Cloud Private Proxy Node	RHEL v7.5
IBM Cloud Private Worker Node	RHEL v7.5
IBM Cloud Private Management Node	RHEL v7.5
IBM Cloud Private VA Node	RHEL v7.5
IBM Cloud Private Boot Node	RHEL v7.5

The software stack relies on server node UEFI settings. Below is a summary of the settings configured for the VMware ESXi hypervisor.

**Table 6. UEFI settings for VMware ESXi**

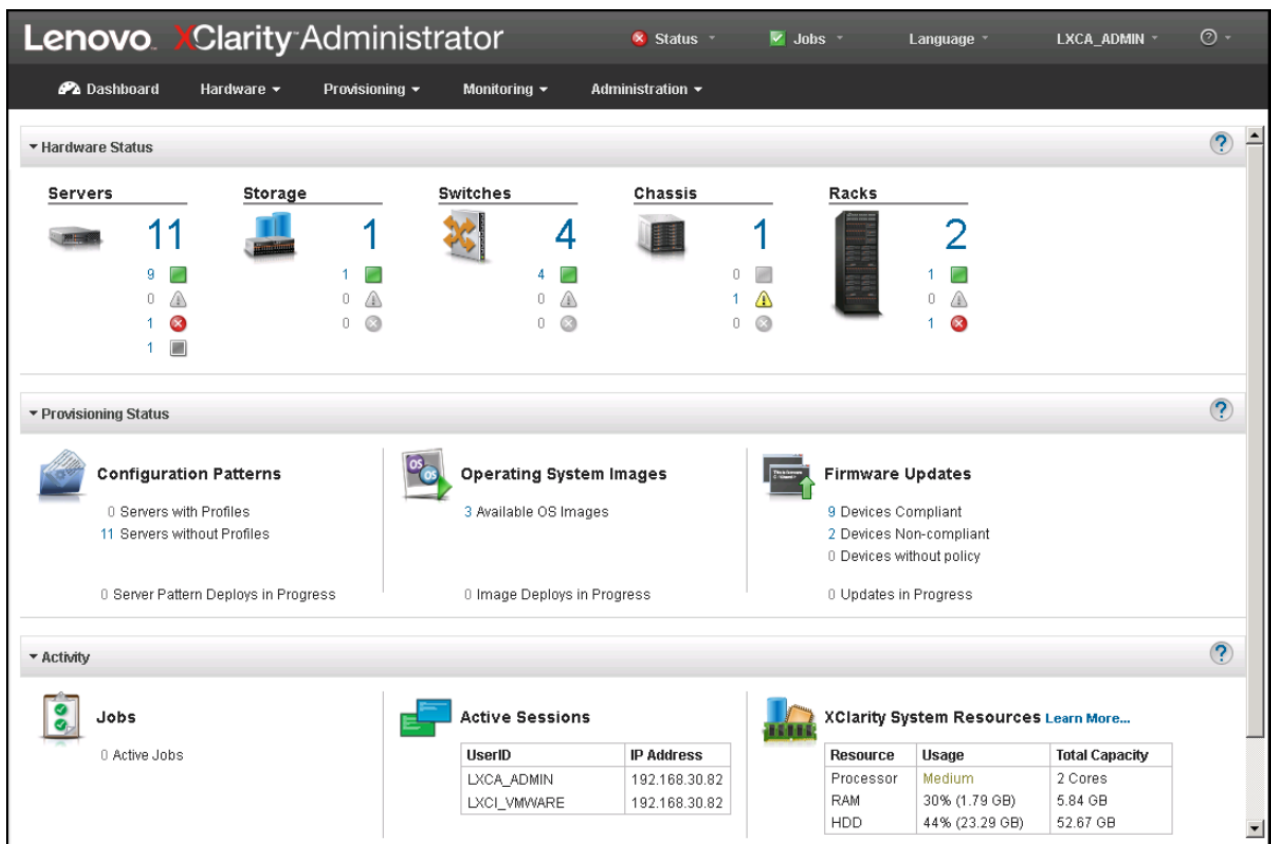
Technologies	Setting
Trusted Platform Module (TPM) 2.0	Enabled
Intel Trusted Execution Technology (Intel TXT)	Enabled
Intel Hyper-Threading Technology (Intel HT Technology)	Enabled
Intel Turbo Boost Technology	Enabled
Intel Speed Shift Technology—hardware P-states (HWP) native	Enabled
Intel Volume Management Device (Intel VMD)	Disabled
Intel Virtualization Technology for Directed I/O (Intel VT-d), virtualization	Enabled
Power-management settings	Performance, workload input/output (I/O) intensive
PCIe, Single-Root I/O Virtualization (SR-IOV) support	Enabled
Integrated I/O	Intel VT-d
Advanced	Set fan profiles to performance PCIe configuration, NIC configuration: disable 3, 4

## 6.4 Systems management

*Systems management* of a cluster includes Operating System and hardware management.

*Hardware management* uses the Lenovo XClarity™ Administrator, which is a centralized resource management solution that reduces complexity, speeds up response and enhances the availability of Lenovo server systems and solutions. XClarity™ is used to install the OS onto new worker nodes; update firmware across the cluster nodes, record hardware alerts and report when repair actions are needed.

Figure 10 shows the Lenovo XClarity™ Administrator interface in which servers, storage, switches and other rack components are managed and status is shown on the dashboard. Lenovo XClarity™ Administrator is a virtual appliance that is quickly imported into a server-virtualized environment.

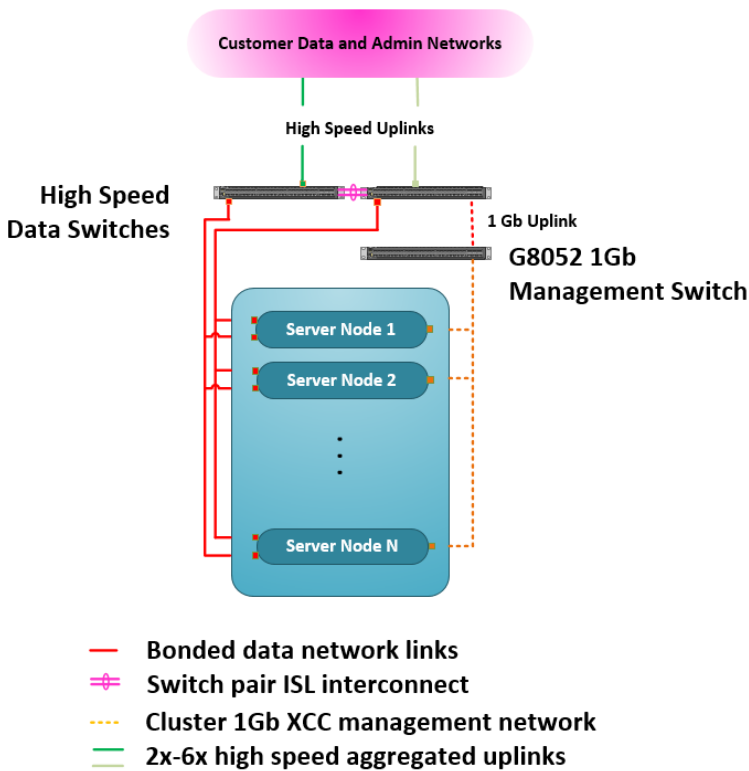


**Figure 10. XClarity™ Administrator interface**

In addition, xCAT provides a scalable distributed computing management and provisioning tool that provides a unified interface for hardware control, discovery and operating system deployment. It can be used to facilitate or automate the management of cluster nodes. For more information about xCAT, see “Resources” section.

## 6.5 Networking

The reference architecture specifies two networks: a high-speed data network and a hardware management network. Two types of rack switches are required; one 1Gb for out-of-band management and a pair of high speed data switches for the data network with High Availability. See figure below.



**Figure 11. Cluster network hardware**

### 6.5.1 Data Network

The data network creates a private cluster among multiple nodes and is used for high-speed data transfer across worker and master nodes, and also for importing data into the cluster. The cluster typically connects to the customer’s corporate data network. This reference architecture demonstrates the Lenovo 10Gb Ethernet System Networking RackSwitch G8272 which provides 48 10Gb Ethernet ports with 40Gb uplink ports. Other available data network speeds are 25Gb with RackSwitch NE2532 or 100Gb with RackSwitch NE10032 using either copper or fiber links.

The two Ethernet NIC ports of each node are link aggregated into a single bonded network connection giving up to double the bandwidth of each individual link. Link redundancy is also provided if one link fails. The two data switches are connected together as a Virtual Link Aggregation Group (vLAG) pair using LACP to provide the switch redundancy. Either high speed data switch can drop out of the network and the other switch continues transferring traffic. The switch pairs are connected with dual 10Gb links called an ISL, which allows maintaining consistency between the two peer switches.

### 6.5.2 Hardware Management network

The hardware management network is a 1 GbE network that is used for in-band operating system administration and out-of-band hardware management. In-band administrative services, such as SSH or Virtual Network Computing (VNC) that is running on the host operating system enables cluster nodes to be administered. Using the XClarity Controller (XCC) within the ThinkSystem SR650 server, out-of-band

management enables the hardware-level management of cluster nodes, such as node deployment or basic input/output system (BIOS) configuration.

Based on customer requirements, the administration links and management links can be segregated onto separate VLANs or subnets. The administrative or management network is typically connected directly to the customer’s administrative network. When the in-band administrative services on the host operating system are used, the cluster is configured to use the data network only.

The reference architecture requires one 1 Gb Ethernet top-of-rack switch for the hardware management network. This rack switch for the hardware management network is connected to each of the nodes in the cluster by using two physical links (one for in-band operating system administration and one link for out-of-band IMM2 hardware management). On the nodes, the administration link connects to port 1 on the integrated 1 GBaseT adapter and the management link connects to the dedicated XCC port.

### 6.5.3 10Gb and 25Gb Data Network Configurations

Both 10Gb and 25Gb high speed data network configurations as well as 100Gb network configurations are possible. The 10Gb network speed is cost effective for most big data workloads. Higher network speeds such as 25Gb and 100Gb data networks are becoming pervasive and increase overall cluster performance and reduce execution time of big data workloads. The table below shows the 10GB components used in this reference architecture and alternate 25Gb components for a higher speed data network.

**Table 7.** 10Gb and 25Gb Network Components

Components	10Gb Network	25Gb Network
High Speed Data Switches	2x, Lenovo RackSwitch G8272,10GbE	2x, ThinkSystem NE2572 RackSwitch, 25GbE
Node network adapter	ThinkSystem Intel X710-DA2 PCIe 10Gb 2-Port SFP+ Ethernet Adapter	ThinkSystem Intel XXV710-DA2 PCIe 25Gb 2-Port SFP28 Ethernet Adapter
Connector type	SFP+	SFP28
Cabling/transceivers	Lenovo Active DAC SFP+	Lenovo 25GBase-SR SFP28 Transceiver

# 7 Deploying VMware vSphere

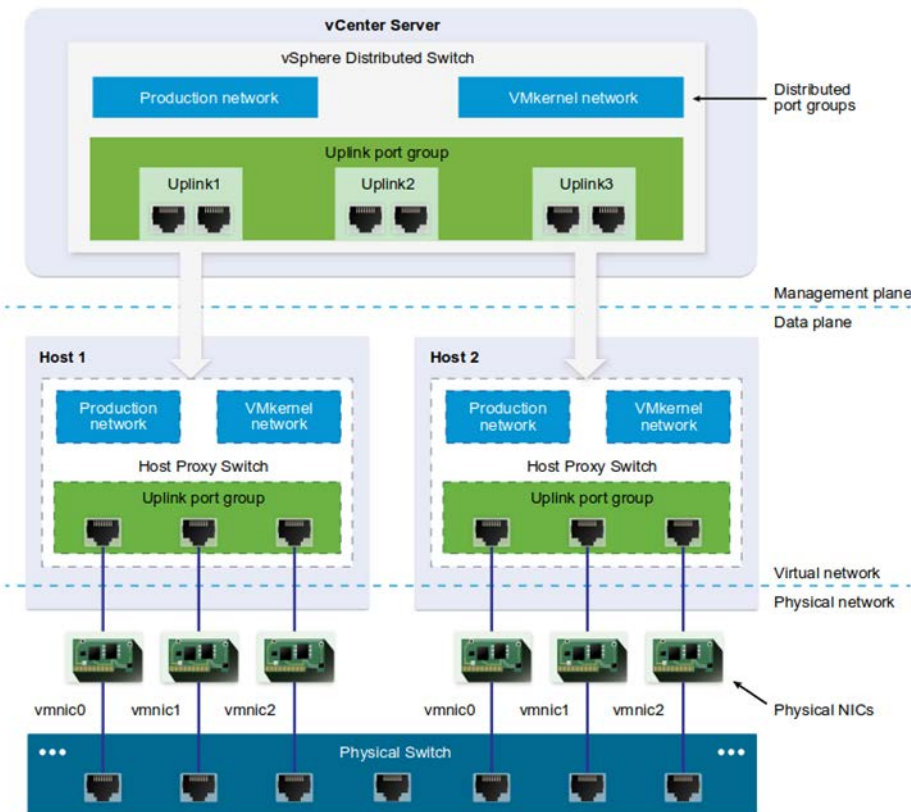
VMware vSphere software includes the ESXi hypervisor for creating Virtual Machines and the vCenter Server for providing a centralized platform for managing the vSphere environment. For this reference architecture the **Embedded Platform Services Controller** option and a **medium VSCA Deployment** size is used and has enabled the ability to bootstrap vSAN on a single host, placing vCenter Server Appliance on the newly created vSAN datastore.

## 7.1 vSphere Configuration

After a successful installation of vCenter, follow these basic steps to finish configuring its deployment:

1. Configure VMware vSphere Distributed Switch networking.

In vSphere, there are two methods of connecting VMs to physical networks. The first is based on vSphere standard switches; the second is based on vSphere Distributed Switch. This reference architecture uses vSphere Distributed Switch. This allows VMs to maintain their network settings even if they are migrated to a different host, as the settings of a distributed switch are propagated to all hosts associated with the switch.



**Figure 12. VMware vSphere distributed switch architecture**

vSphere Distributed Switch separates the data plane and the management plane. Management resides centrally on the vCenter Server and it lets users administer the networking configuration of all

ESXi hosts in a data center. The data plane remains local to every host that is connected to the distributed switch as the host proxy switches. The networking configuration that you create on vCenter Server is automatically pushed down to all host proxy switches.

vSphere Distributed Switch has a lot of other features that distinguish it from a standard vSwitch. One example is the Network I/O Control feature, which allows users to provide quality of service (QoS) on the various traffic types, like the vSAN traffic.

vSphere Distributed Switch specifies two abstractions: Uplink Port Groups and Distributed Port Groups. The first one maps the physical NICs of each ESXi host to uplinks on the distributed switch; the second one provides network connectivity to VMs and forwards the VMkernel traffic. You can configure NIC teaming, failover, load balancing, VLAN, security, traffic shaping, and other policies on distributed port groups.

2. Add and configure VMkernel interfaces for VMware vSphere vMotion and vSAN networks on each ESXi host.

Each ESXi host needs to configure the VMkernel interfaces for vMotion and vSAN networks. Both of these networks are separated and reside on different VLANs. VMware vSphere vMotion provides the functionality of moving VMs automatically across the available pool of resources according to a set of rules, if defined by the administrator.

In vSphere, you can easily create these rules on each ESXi host with the PowerCLI tool, which is the PowerShell interface for managing the whole VMware vSphere stack. Assuming you have created a vSphere Distributed Switch named "Dswitch," then you can use the following commands:

```
New-VirtualPortGroup -Name "vMotion" -VirtualSwitch "DSwitch" -VlanId 11

New-VMHostNetworkAdapter -VMHost 172.17.255.12 -PortGroup "vSAN" -VirtualSwitch
"DSwitch" -IP 192.168.8.12 -SubnetMask 255.255.255.0 -Mtu 9000 -VsanTrafficEnabled $true

New-VirtualPortGroup -Name "vSAN" -VirtualSwitch "DSwitch" -VlanId 12

New-VMHostNetworkAdapter -VMHost 172.17.255.12 -PortGroup "vMotion" -VirtualSwitch
"DSwitch" -IP 192.168.4.12 -SubnetMask 255.255.255.0 -Mtu 9000 -VMotionEnabled $true
```

3. Configure a cluster with vSphere High Availability and VMware vSphere Distributed Resource Scheduler for better HA and scheduling of VMs.

A vSphere Distributed Resource Scheduler cluster is a collection of ESXi hosts and VMs with all shared resources and management interfaces. It offers three cluster-level resource-management capabilities: load balancing, power management, and virtual machine placement. vCenter monitors distribution and usage of CPU and memory resources for all hosts and VMs in the cluster, and then vSphere Distributed Resource Schedulers compares these metrics with the current demand and the imbalance target for the desired resource utilization. vSphere Distributed Resource Scheduler then performs or recommends, depending how is it configured, VM migrations or new VM placement.

vSphere High Availability monitors the pool of ESXi hosts and VMs, and, in the event of host failure, it migrates VMs to hosts with available capacity using vSphere vMotion. When you add new VMs to a vSphere High Availability cluster, vSphere High Availability checks if there are enough resources to power the VM on in case of host failure.

This RA enables HA with host monitoring and vSphere Distributed Resource Schedulers with partially automated settings, where only the placement of VMs onto hosts is automated; the migration recommendations need to be manually applied.

4. Configure vSAN and claim disks for its cache and capacity tier. We recommend that running a health check after IBM Cloud Private is install to ensure that all services are up and running:

```
sudo docker run --net=host -t -e LICENSE=accept -v "$(pwd)":/installer/cluster
ibmcom/icp-inception-amd64:3.1.1-ee healthcheck -vv -f 23
```

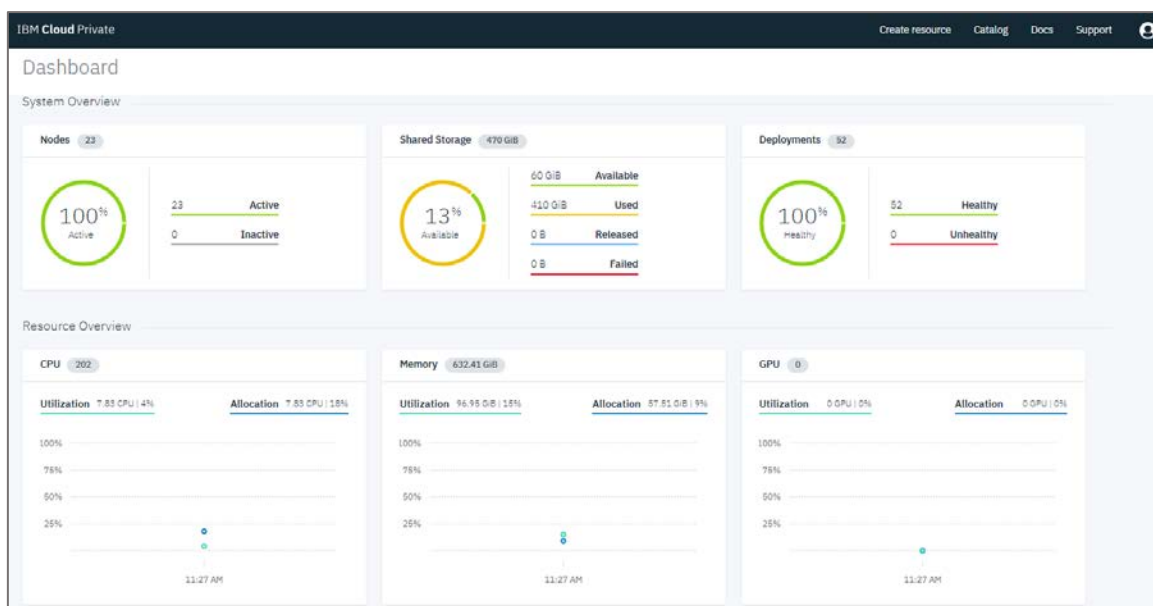


Figure 13. IBM Cloud Private dashboard after a successful deployment and health check

## 7.2 vSphere Distributed Switch

The following section shows configuration of the network environment for the vSphere and IBM Cloud Private cluster.

### 7.2.1 Network Architecture

Table 8. VLAN network addressing

Data Switch VLANs	Server Ports	Example VLAN and IP Range	CIDR	vSphere Distributed Switch
IBM Cloud Private	NIC1, NIC2	Untagged	100.65.80.0/22	DPortGroup-ICP
VMware ESXi (management)	NIC1, NIC2	Untagged	100.65.80.0/22	DPortGroup-MGMT
VMware vSAN	NIC1, NIC2	11	192.168.8.0/24	DPortGroup-vSAN
VMware vSphere vMotion	NIC1, NIC2	12	192.168.12.0/24	DPortGroup-vMotion



**Table 9. vSphere networks using vSphere Distributed Switch**

Data Switch VLANs	Server Ports	Example VLAN	Example IP Range	vSphere Distributed Switch
IBM Cloud Private (VMs)	NIC 1, NIC2	Untagged	172.17.255.0/24	DPort Group: ICP
VMware ESXi (management)	NIC 1, NIC2	Untagged	172.17.255.0/24	DPort Group: MGMT
VMware vSAN	NIC 1, NIC2	11	192.168.8.0/24	DPort Group: vSAN
VMware vSphere vMotion	NIC 1, NIC2	12	192.168.12.0/24	DPort Group: vMotion
<b>Additional Networks</b>				
Hardware network for Intelligent Platform Management Interface (IPMI) management		192.168.24.0/24		
IBM Cloud Private Calico networks on top of the IBM Cloud Private (VMs) network (overlay):				
<ul style="list-style-type: none"> <li>IBM Cloud Private Calico network for pods</li> <li>IBM Cloud Private Calico network for services</li> </ul>		10.1.0.0/16 10.0.0.1/24		

## 7.2.2 Management Network

The management network uses a separate network and cards because it has no dependencies to the other networks in the Intel Reference Solution for IBM Cloud Private. You can use your organizations' own best practices and processes in provisioning it. In this setup, vSAN is configured in an all-flash mode, where one flash device is used for cache, while other flash devices are used for the capacity layer. This greatly improves the performance of VMs running on vSAN, in comparison to a hybrid vSAN configuration. In an all-flash configuration, the cache tier is used for the write cache.

Each ESXi host in the cluster has a configured VMkernel interface for vSAN traffic on separate distributed port groups. For this RA, we created vSAN disk groups by claiming each server's disks; in our case, Intel Optane SSD DC P4800X disks are matched as cache tier and Intel SSD DC P4510 disk as capacity tier, and each disk group contains a maximum of one flash cache device and up to seven capacity disks.

## 7.3 VMware vSAN configuration

Table 10 shows how vSAN is configured for each server in the cluster (SKU plus configuration); there are two disk groups with this configuration.

**Table 10. VMware vSAN configuration**

Disk Group	Storage SKU	Tier	Capacity
Disk Group 1	1 x Intel Optane SSD DC P4800X (2.5-in. PCIe)	Cache tier	350 GB
	3 x Intel SSD DC P4510 (2.5-in. PCIe 3.1)	Capacity tier	5.46 TB
Disk Group 2	1 x Intel Optane SSD DC P4800X (2.5-in. PCIe)	Cache tier	350 GB
	3 x Intel SSD DC P4510 (2.5-in. PCIe 3.1)	Capacity tier	5.46 TB

An all-flash configuration allows users to turn on the deduplication and compression modes, which can help improve the TCO by reducing the data stored on your physical disks.

# 8 Deploying IBM Cloud Private

Installing and configuring IBM Cloud Private is facilitated via Ansible running on the boot node. Details on the specific installation steps can be found at the following link for the version of IBM Cloud Private used in this reference architecture:

[ibm.com/support/knowledgecenter/SSBS6K\\_3.1.1/installing/install\\_containers.html](https://ibm.com/support/knowledgecenter/SSBS6K_3.1.1/installing/install_containers.html).

After IBM Cloud Private is installed it is recommended to run a health check and ensure all services are up and running:

```
sudo docker run --net=host -t -e LICENSE=accept -v "$(pwd)":/installer/cluster ibmcom/icp-inception-amd64:3.1.1-ee healthcheck -vv -f 23
```

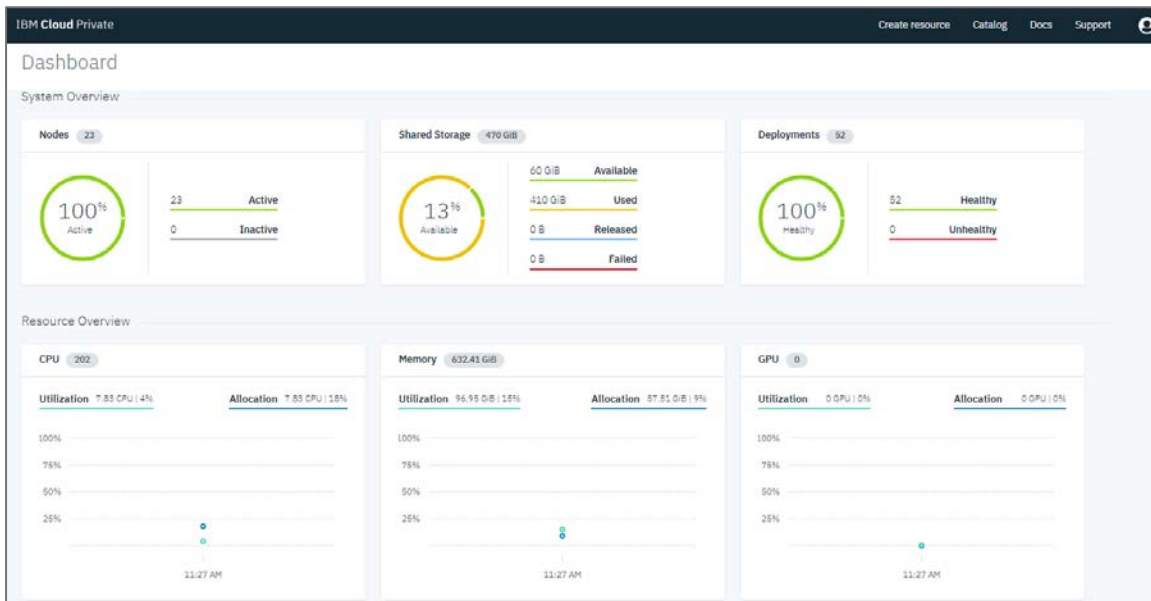


Figure 14. IBM Cloud Private dashboard after a successful deployment and health check

## 8.1 Cluster Monitoring

The IBM Cloud Private software platform provides all necessary tools for monitoring, logging, metering, and alerting. The platform makes it easy to manage processes without specialized expertise.

### 8.1.1 Logging

The ELK stack in IBM Cloud Private combines Elasticsearch, Logstash, and Kibana into one package to use with IBM Cloud Private for logging Kubernetes and Docker events. You can deploy ELK from the IBM Cloud Private catalog to one of the namespaces in the cluster in order for it to log messages from Docker applications.

## 8.1.2 Monitoring and Alerting

IBM Cloud Private uses Prometheus and Grafana to gather and visualize data from the cluster and applications. Prometheus or add-on tools can be used to provide alerts, with or without notifications.

The Grafana dashboard, which is part of the cluster monitoring dashboard provided by IBM Cloud Private, offers fourteen different types of dashboard templates to monitor the status of applications deployed through the cluster. Examples include:

- **Kubernetes Pod Overview**, which contains pod specific metrics as CPU, Memory, Network status and how many restarts single pod has had.



**Figure 15. IBM Cloud Private Grafana Dashboard – Kubernetes POD Overview**

- **Kubernetes Cluster Monitoring**, which provides information collected by Prometheus about the Kubernetes cluster and includes CPU, Memory, and Filesystem usage. Moreover, the dashboard

shows essential information about components, such as individual pods, containers, and systemd services.



Figure 16. IBM Cloud Private Grafana Dashboard – Kubernetes Cluster Monitoring

- **NGINX Ingress Controller**, which includes an array of metrics about the ingress controller, including Requests, Volume, and Network I/O Pressure.

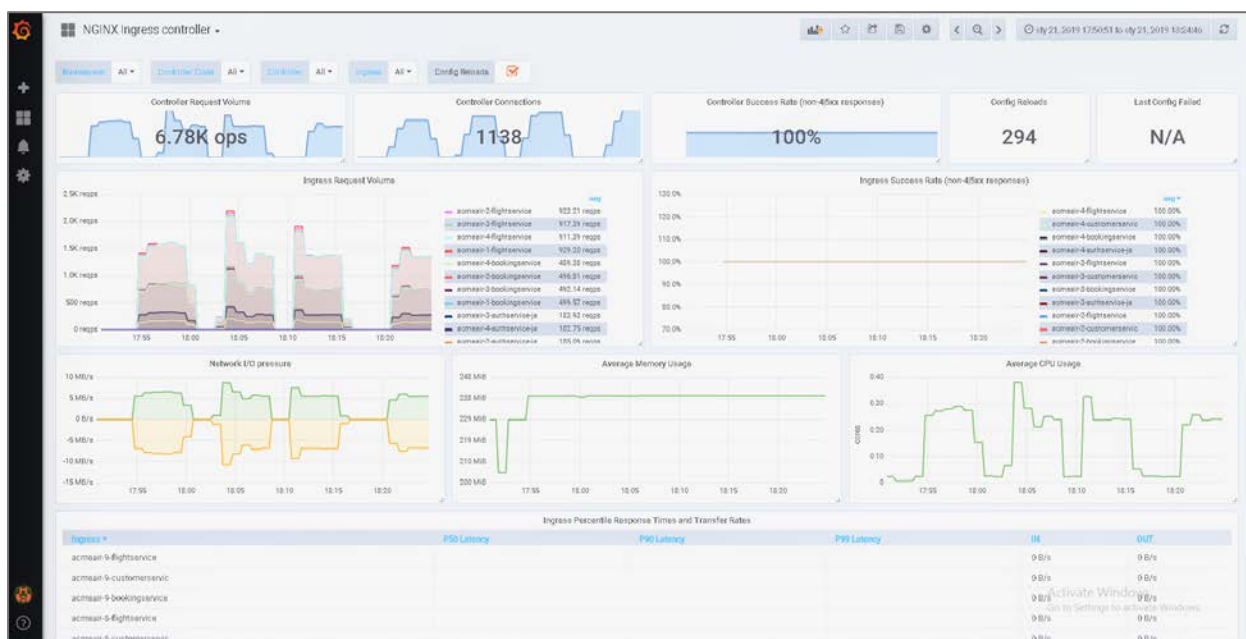
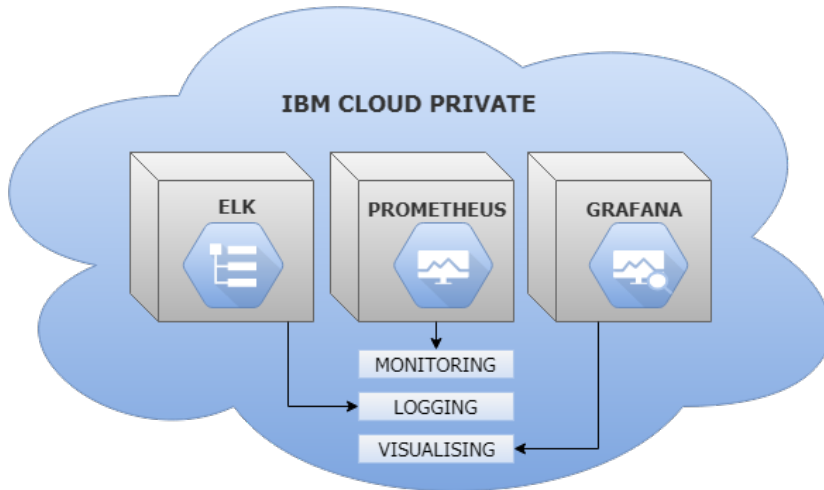


Figure 17. IBM Cloud Private Grafana Dashboard – NGINX Ingress controller

### 8.1.3 Metering

Metering is provided by IBM Cloud Private for tracking application and cluster metrics.



**Figure 18. Logging and monitoring tools in IBM Cloud Private**

## 8.2 Network Configurations - IBM Cloud Private

The default network for IBM Cloud Private pods is provided by Calico, which is configured automatically during the installation process of IBM Cloud Private on top of the network configured for VMs in vSphere.

Calico is a new approach to virtual networking and network security for containers, VMs, and bare-metal services. It provides a rich set of security-enforcement capabilities that run on a highly scalable and efficient virtual network. The Calico node runs in a Docker container on the Kubernetes master node and on each Kubernetes worker node in the cluster. The calico-cni plugin integrates directly with the Kubernetes kubelet process on each node to discover which pods are created and add them to Calico networking.

For this RA's setup, Calico is configured with IP-in-IP encapsulation. You must set this parameter even if all the nodes belong to the same subnet. This configuration enables encapsulation of pod-to-pod traffic over the underlying network infrastructure.

The `calico_tunnel_mtu` parameter must be set based on the maximum transmission unit (MTU) of the interface that is configured for use by Calico.

If the IP-in-IP encapsulation parameter is set to `true`, 20 bytes are used for the IP-IP tunnel header. You must set the `calico_tunnel_mtu` parameter to be at least 20 bytes less than the actual MTU of the interface.

IBM Cloud Private Calico networks on top of the IBM Cloud Private VMs network overlay:

- IBM Cloud Private Calico network for pods, 10.1.0.0/16—The IPv4 network to use for the entire pod network.
- IBM Cloud Private Calico network for services, 10.0.0.1/24—The Kubernetes service cluster IP range. This configuration allocates a block of IPs for services. These service IPs do not need to be routable, since kube-proxy converts service IPs to pod IPs before traffic leaves the node.

The IBM Cloud Private network for pods' underlying network for VMs and the IBM Cloud Private service cluster IP range must not be in conflict with each other.

Three system-configuration parameters on IBM Cloud Private nodes and ingress controller must be set to support a large number of open files and TCP connections with large bursts of messages. Changes can be made using the `/etc/rc.d/rc.local` or `/etc/sysctl.conf` script to preserve changes after reboot.

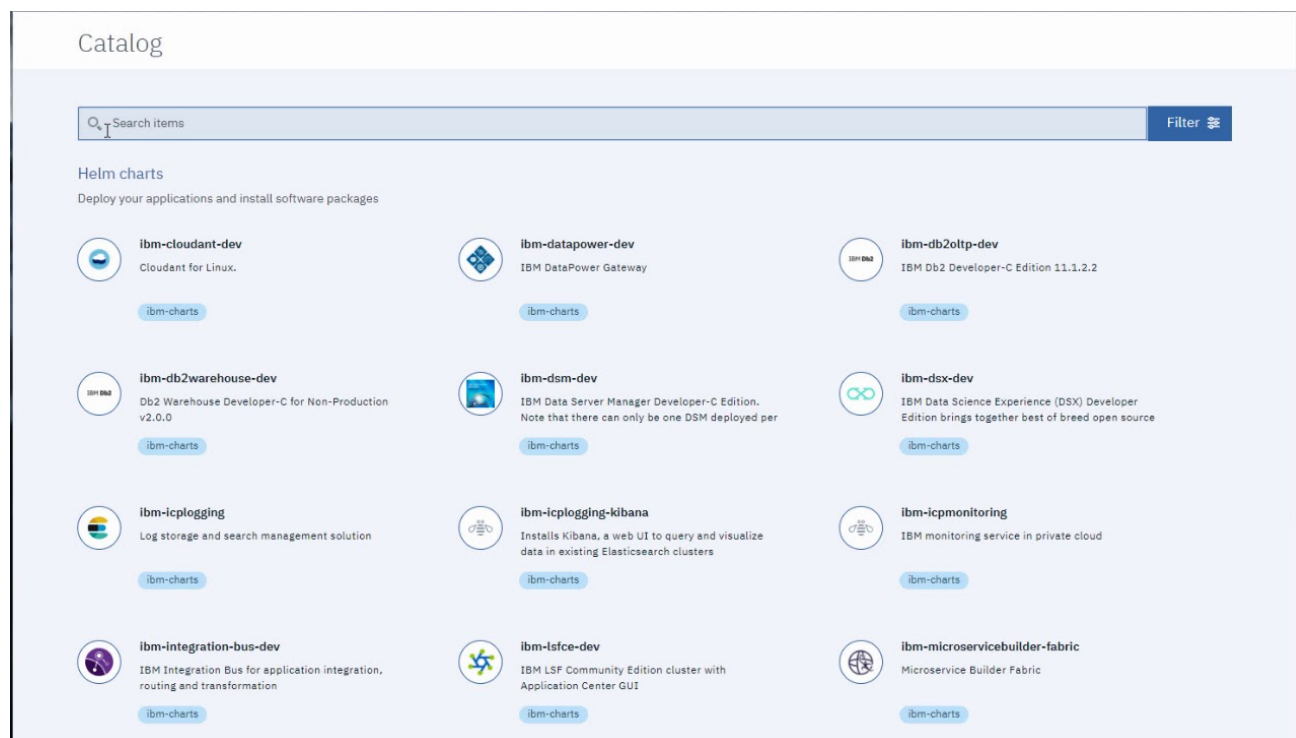
1. `/proc/sys/fs/file-max`: The maximum number of concurrently open files.
2. `/proc/sys/net/ipv4/tcp_max_syn_backlog`: The maximum number of remembered connection requests, which did not receive an acknowledgment from connecting client. The default value is 1,024 for systems with more than 128 MB of memory, and 128 MB for low-memory machines.
3. `/proc/sys/net/core/somaxconn`: Limit of socket `listen()` backlog, known in user space as `SOMAXCONN`. Defaults to 128.
4. `/proc/sys/net/ipv4/ip_local_port_range`: Increase system IP port limits to allow for more connections.
5. `net.ipv4.ip_local_port_range = 1024 65535`

Large numbers of workloads running concurrently may create problems on Linux systems—particularly for proxy nodes—concerning running out of space in the `conntrack` table, which can cause poor iptables performance. To avoid these issues, we increased the number on the proxy nodes:

```
echo 524288 > /proc/sys/net/netfilter/nf_conntrack_max
```

## 9 Deploying IBM Db2 on IBM Cloud Private

Having installed and configured IBM Cloud Private, application packages can be deployed from the IBM Cloud Private catalog. As described above, the Catalog provides a centralized location from which you can browse for and install packages in your cluster (see Figure 19). IBM products such as Db2 Developer-C Edition are included in this catalog to facilitate easy deployment.



**Figure 19. An example of the IBM Cloud Private Catalog screen.**

Details on the specific installation steps can be found at the following link:

<https://developer.ibm.com/recipes/tutorials/db2-integration-into-ibm-cloud-private/>

After installing Db2 Developer-C edition, confirm that the installation is complete using the following commands:

```
kubectl get events
```

```
kubectl get pods; kubectl logs -f <nameofpod>
```

A similar procedure can be used to install other IBM Db2 editions such as the IBM Db2 Direct Advanced Edition and Db2 Advanced Enterprise Server Edition.

### 9.1 High availability considerations

High availability disaster recovery (HADR) provides a high availability solution for both partial and complete



site failures. HADR protects against data loss by replicating data changes from a source database, called the primary database, to the target databases, called the standby databases. HADR supports up to three remote standby servers. More information about HADR capabilities and limitations can be found at the following link:

[https://www.ibm.com/support/knowledgecenter/en/SSEPGG\\_11.1.0/com.ibm.db2.luw.admin.ha.doc/doc/c0011267.html](https://www.ibm.com/support/knowledgecenter/en/SSEPGG_11.1.0/com.ibm.db2.luw.admin.ha.doc/doc/c0011267.html)

# 10 Acknowledgements

---

This reference architecture document has benefited from contributions and careful review comments provided by several colleagues. In particular, we gratefully acknowledge the collaboration and participation by Dan Kangas, Ajay Dholakia, Prasad Venkatachar, Jeff Huxtable and Ron Kunkel from Lenovo; David Solomon from IBM and Małgorzata Rembas, Dominika Krzyszczyk, Karol Sz waj, Filip Skirtun and Raghu Moorthy from Intel.

Reader can contact Ajay Dholakia or Prasad Venkatachar from Lenovo for inquiries and information about the IBM Cloud Private solution described in this reference architecture.

# Resources

---

Lenovo ThinkSystem SR650:

- Lenovo Press product guide: <https://lenovopress.com/lp0644>

Intel Xeon Scalable Family Balanced Memory Configurations Whitepaper

- <https://lenovopress.com/lp0742-intel-xeon-scalable-family-balanced-memory-configurations>

Lenovo RackSwitch G8052 (1GbE Switch):

- Lenovo Press product guide: <https://lenovopress.com/tips1270>

Lenovo RackSwitch G8272 (10GbE Switch):

- Product page: [shop.lenovo.com/us/en/systems/browsebuy/lenovo-rackswitch-g8272.html](http://shop.lenovo.com/us/en/systems/browsebuy/lenovo-rackswitch-g8272.html)
- Lenovo Press product guide: [lenovopress.com/tips1267](https://lenovopress.com/tips1267)

Lenovo ThinkSystem NE10032 (40GbE/100GbE Switch):

- Lenovo Press product guide: <https://lenovopress.com/lp0609>

Lenovo Big Data Validated Design for Hortonworks Data Platform Using ThinkSystem Servers:

- Lenovo Press Solution page: <https://lenovopress.com/lp0776>

Lenovo Big Data Validated Design for IBM SQL Analytics on ThinkSystem Servers:

- Lenovo Press Solution page: <https://lenovopress.com/lp0833>

Lenovo Big Data Validated Design for IBM Db2 Warehouse on ThinkSystem Servers:

- Lenovo Press Solution page: <https://lenovopress.com/lp0846>

Lenovo XClarity Administrator:

- Lenovo Press product guide: <https://lenovopress.com/tips1200>

IBM Cloud Private product page:

- <https://www.ibm.com/us-en/marketplace/ibm-cloud-private/details>

IBM Cloud Private overview:

- [https://www.ibm.com/support/knowledgecenter/en/SSBS6K\\_3.1.1/getting\\_started/introduction.html](https://www.ibm.com/support/knowledgecenter/en/SSBS6K_3.1.1/getting_started/introduction.html)

IBM Db2 Warehouse

- <https://www.ibm.com/products/db2-warehouse>
- <https://www.ibm.com/support/knowledgecenter/en/SS6NHC/com.ibm.swg.im.dashdb.kc.doc/welcome.html>

IBM Spectrum Scale

- [https://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale\\_welcome.html](https://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html)

VMware:

- <https://www.vmware.com/products/vsphere.html>
- <https://www.vmware.com/products/vsan.html>

Intel SSD products:

- <https://www.intel.com/content/www/us/en/products/memory-storage/solid-state-drives/data-center-ssds.html>

# Document history

---

Version 1.0	February 8, 2019	<ul style="list-style-type: none"><li>• First version</li></ul>
-------------	------------------	---

# Trademarks and special notices

---

© Copyright Lenovo 2019.

References in this document to Lenovo products or services do not imply that Lenovo intends to make them available in every country.

Lenovo, the Lenovo logo, ThinkCentre, ThinkVision, ThinkVantage, ThinkPlus and Rescue and Recovery are trademarks of Lenovo.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-Lenovo websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this Lenovo product and use of those websites is at your own risk.