

Release Notes for Virtualized Infrastructure Manager, Version 6.0

First Published Date: December 19, 2025

Contents

- Introduction 1
- Compatibility Versions..... 1
- Hardware Compatibility 1
- Software Compatibility 1
- Features of HVIM 3
- Resolved Bugs..... 11
- Known Caveats 11

Introduction

Network Function Virtualization Infrastructure (NFVI) provides a virtual layer and hardware environment in which virtual network functions (VNFs) operate. VNFs provide a well-defined network function that offers routing, intrusion, detection, Domain Name Service (DNS), caching, Network Address Translation (NAT), and other network functions. While the network functions required a tight integration between a network software and hardware in the past, VNFs decouple the software from the underlying hardware.

NFVI 6.0 is based on the antelope release of OpenStack, an open-source cloud operating system that controls large pools of compute, storage, and networking resources. Our version of OpenStack is Virtualized Infrastructure Manager (VIM) manages the OpenStack compute, network, and storage services, and all NFVI build and control functions.

Key roles of NFVI pods are:

- Control (including Networking)
- Compute.
- Storage
- Management, logging, and monitoring

Compatibility Versions

Hardware Compatibility

Hardware used to create the NFVI pods include:


- Cisco Unified Computing System C240 M5 or C240 M6 or C220 M5 or C220 M6 or C220 M7: Performs management and storage functions, and services. Includes dedicated Ceph (UCS 240-M5 or UCS 240-M6) distributed object store and the file system. Only Red Hat Ceph is supported.
- Cisco UCS C220/240 M5 or M6 or M7: Performs control, compute, and storage services.
- Combination of M5 series servers for VIC or NIC (40G) based Micropod offering.

The UCS C240 and C220 servers are M5/M6 Small Form Factor (SFF) models. For M5, the operating systems boot from Hard Disk Drive (HDD)/Solid State Drive (SDD) for control/compute nodes, and boot from internal SSD for Ceph nodes. For M6, the operating systems boot from Hard Disk Drive (HDD) /Solid State Drive (SDD) for all nodes. HCL supports pure Intel NIC configuration and Cisco 40G VIC with Intel NIC configuration. From this release, Pure Intel NIC configuration is currently supported for M6 and M7.


Software Compatibility

Software applications that manage NFVI hosts and services include:

- Red Hat Enterprise Linux (RHEL) 9.4 with OSP18 rpms and Ceph 7: Provides the core Operating system with OpenStack capability. RHEL 9.4 and OSP 18 are installed on all NFVI UCS servers.
 - HVIM 6.0 is based out of Openstack 18 RPMs only without Openshift.
- VIM: An OpenStack orchestration system that helps to deploy and manage an OpenStack cloud offering from bare metal installation to OpenStack services, considering the hardware and software redundancy, security, and monitoring. OpenStack release with more features and usability enhancements are tested for functionality, scale, and performance.
- Unified Management (UM): Deploys, provisions, and manages VIM on Cisco UCS servers. It provides UI to manage multiple pods, when installed on a dedicated server Unified Management node.
- VIM Monitor: Used to provide integrated monitoring and alerting of the NFV Infrastructure layer.
- Cisco Integrated Management Controller (IMC): When installing VIM, Cisco IMC 4.x or later is recommended.

 Starting from vim 4.4.3, there is no support for vim Fresh Installation on UCS M4 servers.

Following is the list of Cisco IMC versions with which VIM has been tested in the Engineering Lab:

Server	Description
UCS-M4 Servers	Starting from vim 4.4.3, Fresh install on UCS-M4 servers is no longer supported. Kindly plan to migrate/use for UCS M5 or M6 servers.
CS-M5 servers	<p>For UCS-M5-C220 and UCS-M5-240</p> <ul style="list-style-type: none"> • Use the UCS suggested release based on software quality, stability and longevity CIMC 4.3(2.250045)*. HVIM BU tested this version in lab on RHEL VM's. • Do not use 3.1(3c) to 3.1(3h), 3.0(4a), 4.0(2c), or 4.0(2d). <p> There is no support for GPU</p>
UCS-M6 servers	<ul style="list-style-type: none"> • Use the UCS suggested release based on software quality, stability and longevity CIMC 4.3(6.250053)*. CVIM BU tested this version in lab on RHEL VM's

*indicates that it is tested in Intel NIC BOM
 Other versions of CIMC may work, but you must test VIM deployment with those versions before rolling it into production.

We recommend the customers to move forward with the new releases of CIMC versions after testing and validating it in a non-production environment if old releases are deprecated or if new fixes solve their problem.

Layer 2 networking protocols include:

- VLAN supported using Open vSwitch (OVS) must be migrated to OVN
- VLAN support using OVN.
- VLAN supported using ML2/VPP. VPP is currently supported for M6 on Pure Intel NIC pod.

For pods based on C-series with Intel NIC Single Root I/O Virtualization (SRIOV), the SRIOV allows a single physical PCI Express to be shared in a different virtual environment. The SRIOV offers different virtual functions to different virtual components over the same physical NIC.

If you use only VIC, VPP as a mechanism driver is not supported.

Features of HVIM



VIM is the only standalone fully automated cloud lifecycle manager offered for the private cloud. It integrates with Cisco UCS C-series servers, Cisco VIC, and Intel NIC. It helps cloud administrators set up and manage private clouds.

The following are the new features and enhancements as part of the VIM 6.0 release:

Feature Name	Description
Infrastructure, OpenStack & Ceph version	<ul style="list-style-type: none"> • Red Hat 9.4 Kernel with EEUS and OSP 18* <ul style="list-style-type: none"> ○ Kernel Version: 5.14.0-427.88.1.el9_4.x86_64 • Python 3.9.18 • Docker 23.0.2 for management node only. Podman 4.4.1 for cloud nodes • Upgrade from CVIM/HVIM 5.0.x (x=3,4) to HVIM 6.0 • Support for migrating from OVS to OVN. • VPP 25.06 • Kubernetes : 1.29.5
Hardware support matrix	<ul style="list-style-type: none"> • UCS C240/220 M5 controller or compute with Intel Skylake or Cascade Lake UCS C240/220 M6 & C220 M7 controller or compute with Intel Ice Lake • UCS C220/240 M5 in a Micropod environment, with an option to add up to 16 UCS C220/240 M5 computes. • UCS C220 M7 in a Micropod environment* • UCS C240/220 M5 controller or compute with Intel X710 NIC, SR-IOV, and Cisco Nexus 9000 • UCS C240/220 M5 servers with Cisco 1457 (for control plane) and Intel XXV710 NIC (for data plane with VPP) and SR-IOV

Feature Name	Description
	<ul style="list-style-type: none"> • UCS C240/220 M6 servers with Cisco 1467 or 1477 (for control plane) and Intel E810XXVDA2, or E810CQDA2 NIC and SR- IOV • UCS C240/220 M6 servers with Pure Intel NIC E810XXVDA4 BOM for control, data plane, and SRIOV Support of UCS-C240 for edge deployment of the cloud • UCS C220 M7 server with Cisco VIC 15427 (for control plane) and Intel E810XXVDA2 NIC for SRIOV* • SATA M.2 (960G) as an option for a boot drive • Support of both HDD and SSD based management node or UM node. • UCS 220/240 M5 in a nano-pod environment with 2-X710 Intel NIC or 25G/40G VIC/NIC BOM
NIC support	<ul style="list-style-type: none"> • Cisco VIC: VIC 1227, 1240, 1340, 1380, 1387 (for M5) in 40G VIC/NIC offering, 1457, 1467 (M6), 1477 (M6), 15427 (M7)* • Intel NIC: X710, 520, XL710, xxv710 (25G), EA2 (25G for M6), E810CQDA2 (100G for M6), E810XXVDA4 (25G for M6), E810XXVDA2 (25G for M7)
Pod type	<ul style="list-style-type: none"> • Dedicated control, compute, and storage (C-series) node running on Intel X710 (for M5) (full on) with Cisco Nexus 9000 series switch (only for Intel NIC and VPP as mechanism driver) as ToR. Support of UCS-M5 (Cisco VIC 1457 with 2-XL710). • Dedicated control, compute, and storage (C-series) node running on Cisco VIC and Intel NIC (full on) with Cisco Nexus 9000 as ToR. Only SRIOV is supported by Intel NIC. Support of Intel XL710 (2 or 4 NIC cards or compute) on M4/M5 pods for SRIOV cards in the VIC/NIC combination. • Support of Intel E810XXVDA2 (2 cards) or Intel E810CQDA2 (2 cards) on M6 pods for SRIOV cards in the VIC/NIC combination. • Support of Intel E810XXVDA2 (2 cards, 2*25G on each card) on M7 pods for SRIOV cards in the VIC/NIC combination. • Support of Intel E810XXVDA4 (2 cards, 4*25G) on M6 pods for samx,pet and SRIOV in the pure Intel NIC combination Dedicated control, compute, and storage (UCS M5 SFF C-series) node running on Cisco VIC 1457 and Intel xxv710 NIC (full on) with Cisco Nexus 9000 as ToR. Only SRIOV is supported by Intel NIC. With VPP and OVS as the mechanism driver, the number of SRIOV ports are 2 or 4, respectively. • Dedicated control, compute, and storage (UCS M6 SFF C-series) node running on Cisco VIC 1467 and Intel E810XXVDA2 NIC (full on) with Cisco Nexus 9000 as ToR. Only SRIOV is supported by Intel NIC. • Dedicated control, compute, and storage (UCS M6 SFF C-series) node running on Cisco VIC 1477 and Intel E810CQDA2 NIC (full on) with Cisco Nexus 9000 as ToR. Only SRIOV is supported by Intel NIC. • Micropod: Integrated (AIO) control, compute, and storage (C-series) node running on Cisco VIC, Intel X710X, or VIC and NIC combinations. Micropod


Feature Name	Description
	<p>can be optionally expanded to accommodate more computes (up to 16) running with the same NIC type. This can be done as a Day 0 or Day 1 activity. The computes can boot off HDD or SSD. From VIM 3.4.1, the Micropod option has been extended to Quanta (D52BE-2U) servers with Intel XXV710 NIC (25G) and Cisco Nexus 9000 (-FX series) as ToR.</p> <ul style="list-style-type: none"> • In a full-on (VIC based), or hyper-converged pod, computes can either have a combination of 1-Cisco VIC (1227) and (2x10GE 520/2x40GE 710XL Intel NIC) or 1-CiscoVIC (1227). The compute running pure Cisco VIC does not run SR-IOV. HCL supports HP DL360 Gen9 and Dell PowerEdge R740 as a third-party compute. • VIM does not support a combination of computes from different vendors. • On a Full-on pod, mix of Cisco UCS M5 and Cisco UCS M6 for storage nodes
Tor Support	<ul style="list-style-type: none"> • NXOS support (preferably N9K)
IPv6 support for management network	<ul style="list-style-type: none"> • Static IPv6 management assignment for servers. • Support of IPv6 for NTP, DNS, LDAP, external syslog server, and AD. Support of IPv6 for the cloud API endpoint. • Support of CIMC over IPv6. RestAPI over IPv6. • Support for IPv6 filters for administration source networks. Support of UM over IPv6
Mechanism drivers	<ul style="list-style-type: none"> • OVS/VLAN (Deprecated in OSP18) • VPP (25.06)/VLAN (Fast Networking, Fast Data VPP/VLAN, based on the VPP fast virtual switch over intel NIC) • OVN/VLAN* • OVS to OVN Migration*
Installation or update method	<ul style="list-style-type: none"> • Fully automated online or offline installation. • Support of offline installation via USB or file-based image. • Support of Cisco VIM Software Hub to mitigate the problem associated with the logistics of USB distribution for air-gapped installation. • Support of USB 3.0 128GB for M5, M6, M7.

Feature Name	Description
Scale	<ul style="list-style-type: none"> Full Pod: Total of 128 nodes (compute and OSD) with Ceph OSD max at 20. <div style="border: 1px solid #f0e68c; padding: 5px; margin: 10px 0;">  The 128-node limit is based on the scale testing done with OVS and based on customer deployments. </div> <p>Ensure that you deploy a maximum of 60 nodes at a time. Also, after Day 0, you can add only/remove one ceph node at a time.</p> <p>The 32-node limit is based on the scale testing done with VPP.</p> <ul style="list-style-type: none"> Micro/Edge pod: Supports a maximum of 16 standalone compute nodes. <div style="border: 1px solid #f0e68c; padding: 5px; margin: 10px 0;">  Ceph OSDs can be HDD or SSD based but must be uniform across the pod. Computes can boot off 2x1.2TB HDD or 2x960 GB SSD). In the same pod, some computes have SSD, while others can have HDD. </div> <p>Contact the VIM product management team for specific use case and BOM details applicable for each type of pod.</p>
Automated pod life cycle management	<ul style="list-style-type: none"> Add or remove compute and Ceph nodes and replace the controller node. Static IP management for storage network. Reduction of tenant or provider VLAN through reconfiguration to a minimum two. Reconfiguration of passwords and selected optional services. Automated software update.
Platform security	<ul style="list-style-type: none"> Secure OS, RBAC, network isolation, TLS, source IP filtering (v4 and v6), Keystone v3, Bandit, CSDL-compliant, hardened OS, and SELinux. Change CIMC password post-installation for maintenance and security. Non-root login for administrators. Read-only role is available for OpenStack users. Enabling custom policy for VNF Manager. Option to disable the management node reachability to the cloud API network. Hosting of Horizon behind NAT or with a DNS alias. Cinder volume encryption using Linux Unified Key Setup (LUKS). Support of configurable login banner for SSH sessions. Access to management node using Open LDAP or MS AD on both secure and non-secure. Support for IPv6 filters for administration source networks. Introduction of Vault to encrypt secrets with reconfigure option. Enablement of Vault as an option on Day 2. Extended <i>permit_root_login</i> to Unified Management node. CIMC authentication using LDAP. Support of out of band and inband RedHat identity, policy, and audit (IPA) system.

Feature Name	Description
	<ul style="list-style-type: none"> • Support of Horizon and Keystone login settings. • Support of LDAP on Unified Management node using Open LDAP or MS AD on both secure and non-secure. • SSH and password vulnerabilities for management node. • Kernel changes to address vulnerabilities. • FQDN support for VIM management API. • LDAP support for Kibana. • Support of TLS over management network. • Support of Barbican using HSM. • Ability to disable VM live migration, snapshot, crash dump, and console. • VIM advanced OpenStack policy for operator role • Disable IPA Cache Credentials • Lock Docker Registry when not used in Management Node • OpenStack services container not to start with standard root account except for nova-libvirt • RESTAPI DB password in VAULT. • Securing the cinder used with NAS (Network Attached Storage) • Support of LDAP on SDS node using Open LDAP or MS AD on both secure and non-secure • Support of LDAP on VIM-MON HA Management node using Open LDAP or MS AD on both secure and non-secure • Removing the WebDAV from Apache installations • Improving Network Security with sysctl variables • Password encryption in VIM setup_data. yaml • set the umask value to 027. • Support for TLSV3 • Securing Memcached • TOR password encryption in setup_data
Enhanced Platform Awareness (EPA)	<ul style="list-style-type: none"> • Supports NUMA, CPU pinning, huge pages, and SRIOV with Intel NIC. • Ability to set HYPERTHREADING at a global level or per compute basis. • Ability to set VM_HUGEPAGE_SIZE and VM_HUGEPAGE_PERCENTAGE at a global level or per compute basis. • Ability to allocate user-defined CPU (up to 6) cores to VPP. • Ability to bring in trusted_vf as a reconfigure option on a per server basis. • Ability to allocate user-defined CPU (up to 12) cores to Ceph for Micropod and hyper-converged nodes. Ability to allocate user-defined CPU (up to 30) cores to compute nodes. • Ability to set LIBVIRT_WRITETHROUGH_CACHE on a global or per compute basis. Ability to set ISOLCPU at a global level or per compute basis. • Ability to set VIC_link_training at global level or per compute basis. Improve OSD failure detection time.

Feature Name	Description
	<ul style="list-style-type: none"> • Ability to allocate user-defined memory (16 to 32 GB) to Ceph for Micropod nodes. • Ability to allocate user-defined memory (25 to 500 GB) for compute nodes. • Ability to support custom mapping between OpenStack physnet to SRIOV PF • Option to configure the queue size per VF on sriov only for M6 Server with ICE Driver *
HA and Reliability	<ul style="list-style-type: none"> • Redundancy at hardware and software level. • Automated backup and restore of the management node. Encryption of management node backup snapshot • Relaxation of Security Enhanced Linux (SELinux) requirement for backup and restore of the management node. • Optimization of automated backup and restore of the management node in a connected installation.
Unified Management (UM) support	<ul style="list-style-type: none"> • Single pane of glass in a single mode. Supports multi-tenancy and manages multiple pods from one instance. • EFK Log Management for Insight pod • EFK Log rotation support for insight pod
VIM monitor	<ul style="list-style-type: none"> • Collect the metrics from the entire pod. Supports customizing alerts, sending SNMP traps, and exporting to external metric collectors.
Central logging	<ul style="list-style-type: none"> • EFK integrated with external Syslog (over v4 or v6) for a log offload, with optional support of NFS with EFK snapshot. • EFK over v4 support for VIM Insight pod
External Syslog servers	<ul style="list-style-type: none"> • Supports multiple external Syslog servers over IPv4 or IPv6. The minimum and maximum number of supported external Syslog servers are 1 and 3, respectively.
VM migration	<ul style="list-style-type: none"> • Cold migration and resizing. • NUMA-aware live migration with Virtio and SRIOV.
Pod Migration	<ul style="list-style-type: none"> • Seamless migration of a Micropod to a full pod.
Storage	<ul style="list-style-type: none"> • Block storage with Ceph or NetApp. • Option to have multi-backend (HDD and SSD based) Ceph in the same cluster to support various I/O requirements and latency. • Support for multiple NFS Shares • Enable multiple Netapp backends for Cinder

Feature Name	Description
Monitoring	<ul style="list-style-type: none"> • HA VIM MON initial bringup with rhel9.4/K8s 1.29.5. • Monitor the VIM pods centrally using the Highly Available VIM Monitor (HA VIM-MON) over v4 and v6. • Monitor the VIM pods individually using the local VIM Monitor (VIM-MON) over v4 and v6. • VIM MON local LDAP supports Grafana, Prometheus, and Alert Manager. • Support of non-VIM managed external servers running RHEL or CentOS. • Ceilometer for resource tracking and alarming capabilities across core OpenStack components are applicable only for fullpod. • EFK for VIM-MON HA Pod • Reconfigure support on EFK stack components on HA VIMMON OSD number to differentiate alarms. • VIM-mon local prometheus metrics to external grafana cloud • Disk failure prediction.
Optional OpenStack features	<ul style="list-style-type: none"> • Enable trusted virtual function on a per-server basis. • DHCP reservation for virtual MAC addresses. • Enable VM_HUGE_PAGE_SIZE and VM_HUGE_PAGE percentage on a per-server basis. • Enable CPU and RAM allocation ratio on a per-server basis via add/remove compute or reconfigure. • Flatten volume from snapshot as an option in cinder. • CEPH PG autoscaling on day-0 and option to enable in day-1 via reconfigure. • Enable logging of keystone user logins
Support of external authentication system	<ul style="list-style-type: none"> • DAP with anonymous bind option and • Support for Active Directory (AD).
Software update	<ul style="list-style-type: none"> • Update of cloud software for bug fixes on the same release. • Support for VMs associated with more than 8vCPUs with MultiQ enabled Metadata and VPP plugin.
CIMC upgrade capability	Central management tool to upgrade the CIMC bundle image of one or more servers.
VPP port mirroring	Ability to trace or capture packets for debugging and other administrative purposes. Automated update of BMC or BIOS and firmware of Quanta server.

Feature Name	Description
VXLAN extension into the cloud	<ul style="list-style-type: none"> Extended native external VXLAN network into VNFs in the cloud. Support of Layer 3 adjacency for BGP. Support of single VXLAN network or multi-VXLAN network (with head-end-replication option) terminating on the same compute node. Support of re-binding of Neutron port to another port. Support of L3 fabric via VXLAN. <div style="border: 1px solid orange; padding: 5px; margin-top: 10px;">  Only two-VXLAN network is supported. </div>
Technical support for CIMC	Collection of technical support for CIMC.
Enable TTY logging as an option	Enables TTY logging and forwards the log to an external syslog server and EFK stack running on the management node.
Unified Management authentication	Authentication support through local and LDAP.
CIMC authentication using LDAP	Authentication support through LDAP.
Automated enablement of Intel X710/XL710 NIC's PXE configuration on Cisco UCS-C series	Utility to update Intel X710/XL710 NIC's PXE configuration on Cisco UCS C-series.
Power management of computes	Option to selectively turn OFF or ON the power of computes to conserve energy.
Fan policy for servers	Option to set fan policy globally for UCS C-series based pod.
Disk maintenance for pod nodes	Ability to replace faulty disks on the pod nodes without the need to add, remove, or replace node operation.
Branding of VM workload	Ability to check whether the VMs are running on VIM platform.

Feature Name	Description
Cloud adaptation for low latency workload	<ul style="list-style-type: none"> • Enable real-time kernel to support edge pod. Automated BIOS configuration. • Introduction of custom flavor. • Support of Intel N3000 card on selected servers to handle vRAN workloads. Support of Cache Allocation Technology (CAT) to handle vRAN workloads. • Support of INTEL_SRIOV_VFS (SRIOV support for Intel NIC) and INTEL_FPGA_VFS (support for Intel N3000 FPGA card) at a per-server level
Integrated test tools	<ul style="list-style-type: none"> • Open-source data plane performance benchmarking: VMTP (an open-source data plane VM to VM performance benchmarking tool) and NFVbench (NFVI data plane and a service chain performance benchmarking tool). • Extending VMTP to support v6 over the provider Network. NFVbench support for VXLAN. • Services Health Checks Integration: Cloudpulse and Cloudsanity. • NFVbench support on Intel E810 cards
Optional Settings for Cisco UCS-C Series servers (M4/M5/M6)	<ul style="list-style-type: none"> • Power restore policy. • Raid controller cache write policy Raid level setup for computes
Additional Management Node Settings	<ul style="list-style-type: none"> • Configurable Fluentd log forward delay
HVIM Upgrade	<ul style="list-style-type: none"> • Upgrade support from CVIM/HVIM 5.0.X (X=3,4) to HVIM 6.0*

*Indicates the features introduced in VIM 6.0

- For supported BOM details, reach out to nfvi-plm@hcltech.com or your Account Team.
- Configure LACP on the data plane ports of the Cisco Nexus 9000 ToR when vim is running on Intel NIC for data plane with VPP as the mechanism driver.

Known Caveats

The following features are not supported in VIM 6.0, but are planned for a subsequent release:

- Hvim osdmgmt for replace-osd, replace-journal

- Hvim expand-storage
- Telegraf not reporting SRIOV port metrics
- Srio_queues_per_vf not supported for intel nic 710 series cards.
- Upgrading elastic search to latest version

The following features are no longer supported -- ACI Plugin Integration, Ironic Support, Ceilometer, SolidFire, Swift, NCS-5500 as TOR, Hyperconverged Pods, Third party servers like HP, DELL, Quanta.

Starting from hvim 5.0.0 the following features are no longer supported -- Pod type Edge, Pod type Ceph (central ceph), Pod Type UMHC, RT Kernel, Zadara storage backend, GPU, vGPU, NFVIMON.

To search for a specific bug or to search for all known issues in this release, please contact HCLTECH L1 support at cs-vim@hcltech.com.

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE LICENSE AGREEMENT SIGNED BETWEEN THE PARTIES.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL

FAULTS. HCLSOFTWARE AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL HCLSOFTWARE OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF HCLSOFTWARE OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

HCLSOFTWARE Logo is trademark of HCL Technologies Ltd., and/or its affiliates in the U.S. and other countries.

Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between HCLSOFTWARE and any other company.

Copyrights @HCL Technologies Ltd., 2025 All rights reserved.