

IBM



Software Defined Network for Virtual Environments

Version 1.2, VMware Edition

User Guide

First Edition (August 2014)

© Copyright IBM Corporation 2014

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Preface	11
Who Should Use This Guide	11
What You'll Find in This Guide	11
Typographic Conventions	14
How to Get Help	15
 Part 1: Getting Started	 17
 Chapter 1. IBM SDN VE Introduction	 19
IBM SDN VE Solution Overview	19
Product Editions	20
IBM SDN VE Solution Components	21
Unified Controller	21
IBM SDN VE Additional Features	22
IBM SDN VE Solution Benefits	23
No Disruption to Existing IPv4 Networks	23
Extending the VM Strategies into the Network.	24
High Availability	24
Connectivity Service	25
Unified Controller	25
Enhanced Multi-tenancy for Cloud Providers.	25
Datacenter Consolidation.	25
Maximizing Servers	25
Optimizing Provisioning with Programmable APIs	26
IBM SDN VE Solution Elements	26
Prerequisites	27
Installation Summary	28
Configuration Summary	29
 Chapter 2. Installing Unified Controller Modules	 31
Deploying the Unified Controller Software	31
Install Unified Controller on Host	31
KVM Environment	31
VMware Environment.	34
Initial Unified Controller Setup	40
Start the Unified Controller Module	40
Set the Language	41
Set the Unified Controller IPv4 Address (Optional)	41
Enter License Information	42
Establish SDN VE Controller HA	43
Log in to the Unified Controller	45
Enter Configuration Mode	45
Establish Unified Controller High-Availability	45
Configure SDN VE DOVE HA	45
The Graphical User Interface	46
Next Steps	47

Chapter 3. Installing DSA Modules	49
Deploying the DSA Software	49
Install Using OVF Tool	50
Install Using VMware vSphere Client	50
Initial DSA Setup	55
Start the DSA Module	56
Log In to the DSA	56
Enter Global Configuration Mode	56
Configure the DSA IPv4 Address (Optional)	56
Attach to the SDN VE Controller Cluster IPv4 Address	57
Specify DSA Roles	58
Configure Tunnel Endpoints	59
Next Steps	60
 Chapter 4. Installing the SDN VE 5000V Distributed vSwitch	 61
Deploying the 5000V Controller Software	61
Install Using OVF Tool	61
Install Using VMware vSphere Client	62
Initial 5000V Controller Setup	67
Start the 5000V Controller	67
Set the language	67
Examine the License Agreement	68
Log In to the 5000V Controller	68
Enter Global Configuration Mode	68
Verify the 5000V Controller Version	68
Configure the 5000V IPv4 Addresses (Optional)	68
Create the Global vDS Instance	69
Attach to the DMC Module Cluster IPv4 Address	70
Next Steps	70
 Chapter 5. Virtual Network Configuration	 71
Overview	71
Overlay Configuration	72
Create Tenants	72
Create Connectivity Groups	72
Create Subnets	73
Bind Subnets to the Connectivity Group	74
Define Connectivity Group Policy (Optional)	74
Export Networks to the SDN VE 5000V vSwitch	74
Externalizing the Overlay Networks	75
Configure a VLAN Gateway	76
Configure an External Gateway	76
Configuration of Gateway Interfaces	78
5000V Host Module	79
Install 5000V Host Module	79
Preconditions	79
Copy 5000V vDS Host Module File to ESXi Machines	81
Install 5000V vDS Host Module VIB	81
Configure the Underlay (Physical) Networks at the Unified Controller	81
Attach ESXi Hosts to vDS	83
Configure TEPs	84
Attach End Systems	87

Chapter 6. Network Services	89
Logical Groups	89
Creating a Tenant	89
Creating a Logical Group	90
Creating a Subnet	92
Create a Port	93
Assign Subnet to a Connectivity Group	94
Layer 3 Configuration	95
Connectivity Group Policy	96
Adding a Policy Between Two Connectivity Groups	96
Monitor/Redirect Sessions	97
Create a Replication/Redirection Session	98
Start/Stop/Delete a Monitor/Redirect Session	99
Static Flows	99
Create a Static Flow Set	99
Create a Static Flow	100
Install a Static Flow Set	102
Delete a Static Flow or Uninstall a Flow Set	102
Chapter 7. Topology	103
Topology Manager	103
Actions	103
Search	104
Logical Groups	104
Viewing the Logical Group	104
Viewing Logical Group Properties	104
Physical Networks	105
Viewing the Physical Network	105
Viewing Properties of the Physical Network	105
Viewing the Connectivity Tree	106
Viewing the Flows for a Switch	106
Chapter 8. System Administration	107
User Management	107
Creating users	107
Editing users	109
Deleting users	110
Resetting Passwords	110
Changing Passwords	111
Logging out of IBM Unified Controller GUI	112
Save Configuration	112
System Commands	112
SDN VE HA Cluster Management	114
Rejoining a Cluster	114
Disconnecting from a Cluster	114
Log Settings	115

Remote Server Setup (LDAP / RADIUS)	116
Configuring RADIUS Server	118
Managing LDAP Server	119
Adding LDAP Server.	119
Modifying Domain Name	120
Enabling/Disabling LDAP Service.	120
Deleting LDAP Configuration	121
Managing RADIUS Server	121
Adding RADIUS Server.	121
Modifying Password	122
Enabling/Disabling RADIUS Service.	122
Deleting RADIUS Configuration	122
Managing Configuration.	123
Backup Configuration	123
Restore Configuration	124
Part 2: Advanced Features	125
Chapter 9. OpenStack	127
OpenStack Integration with SDN VE Plugin	128
Plugin Integration	128
Chapter 10. Waypoint Connectivity Service	135
Waypoint Service Operation	135
Transparent Mode	135
Routed NAT Mode	135
Routed Mode	136
Routed Explicit Devices	136
Routed Implicit Devices	136
Waypoint Connectivity	136
Waypoint Discovery	137
Waypoint Configuration	137
Waypoint Configuration Using Service Templates and REST APIs	138
Defining a Connectivity Instance	139
Service Template Example	140
Waypoint Configuration Using Controller GUI.	141
Providing Middlebox Specifications	141
Configuring A Service Chain.	142
Defining a Policy	143
External Connectivity Groups - SNAT Pool Configuration	144
Waypoint High-Availability/Load Balancing	145
Transparent Mode	145
Routed/Routed NAT Mode	145
Limitations	145
Chapter 11. NIST	147
Enabling NIST	147
Acceptable Cipher Suites	148
LDAP Configuration.	149

Chapter 12. Public Key Infrastructure.	.151
PKI Configuration	.151
DSA Configuration	.153
5000V Agent Host Configuration	.153
5000V Controller Configuration	.154
Deleting Certificates	.154
Authentication	.154
Enabling Authentication on the DSA	.156
Enable Authentication on the 5000V Controller	.156
IP Security	.156
Chapter 13. Access Control Lists	.157
MAC Extended ACLs	.157
IPv4 ACLs	.157
Summary of Packet Classifiers	.158
Summary of ACL Actions	.160
Assigning Individual ACLs to a Port	.161
Assigning Individual ACLs to a VNIC Profile	.161
Viewing ACL Statistics	.161
Deleting ACLs	.162
ACL Configuration Examples	.162
Chapter 14. Quality of Service.	.167
QoS Overview	.167
Using DSCP Values to Provide QoS	.167
Differentiated Services Concepts	.167
QoS Levels	.168
Using 802.1p Priority to Provide QoS	.168
QoS Implementation	.169
Rate Limiting	.169
DOVE Connectivity Group	.170
Limitations	.170
Chapter 15. sFlow.	.171
Enabling sFlow	.171
Global Packet Sampling	.171
Statistical Counters	.172
Custom Sampling Groups	.173
sFlow Configuration Information	.175
sFlow Configuration Example	.175
Firewall Considerations	.176
Chapter 16. TCP Segmentation Offload	.179
VXLAN Port	.179
Chapter 17. Virtual Router Redundancy Protocol	.181
VRRP Overview	.181
VRRP Components	.181
Selecting the Master VRRP Router	.181
VRRP Implementation	.182
Configuring VRRP	.182

Part 3: Command Reference 185

Chapter 18. Command Basics 187

Login	187
Command Modes	188
Unified Controller.	189
Privilege EXEC Mode	189
Global Configuration Mode	189
SDN VE DOVE Controller	189
SDN VE DOVE Configuration Mode.	190
Global Commands.	190
Idle Timeout	191

Chapter 19. Show Commands 193

Cluster Information.	193
SPARTA Information	193
Flow Information	194
Connectivity Group Information	198
Host Information.	201
LDAP Server Information	201
Log Information	202
Log Levels	202
View Logs	203
Multicast Information	204
OpenFlow Information	205
Port Information	206
RADIUS Server Information	207
Replication Session Information	208
Subnet Information.	209
Switch Information	209
System Information	210
Running Configuration Information	218
Tenant Information.	219
Topology Information.	219
System Upgrade Information.	220
Users Information	221
SDN VE Version Information	221
Connectivity Group Policy Information	222
NIST Information	222
DOVE Configuration Information	223

Chapter 20. Configuration Commands	.231
Global Configuration Mode	.231
Cluster Configuration	.231
Flowset Configuration	.232
LDAP Server Configuration	.233
ldap server domain	.234
ldap server primary	.234
no ldap server primary	.235
Log Setting Configuration	.235
Multicast Configuration	.237
NIST Configuration	.238
Pagination Configuration	.238
RADIUS Server Configuration	.238
Reset User Password	.240
SDN VE Configuration	.240
Switch Configuration	.240
Tenant Configuration	.241
User Configuration	.243
System Configuration Commands	.245
system ipmgmt ip	.247
system ipmgmt nameserver	.247
system ipmgmt nexthop	.248
no ipmgmt nameserver	.248
no ipmgmt nexthop	.248
system sdn-ve log rm	.250
system sdn-ve authenticate	.251
Log Level	.252
Syslog Enable or Disable	.252
Console Log	.253
Flow Group Configuration Mode Commands	.255
Tenant Configuration Mode Commands	.258
Connectivity Group Configuration	.258
Connectivity Group Policy Configuration	.261
Flow Replication Configuration	.263
Flow Redirection Configuration	.266
Subnet Configuration	.269
Group Configuration Mode Commands	.271
SDN VE DOVE Configuration Mode Commands	.274
Service Gateway Configuration Mode Commands	.284
Miscellaneous Commands	.287
Chapter 21. DSA Show Commands	.289
Chapter 22. DSA Configuration Commands	.297
Clear Commands	.298
CLI Timeout Commands	.299
Controller Commands	.300
Hostname Commands	.301
Image Upgrade Commands	.302
IP Management Commands	.303
Password Configuration Commands	.305
Miscellaneous Commands	.306
PKI Configuration Commands	.308
Terminal Length Configuration Commands	.311

Chapter 23. Diagnostics Commands	313
Part 4: Appendices	317
<hr/>	
Appendix A. New and Updated Features	319
CLI	319
DOVECLI.	319
High-Availability (HA).	319
NIST	320
Overlay Networks	320
PKI.	320
QoS	320
TSO	320
VRRP.	321
Waypoint Connectivity Service	321
Appendix B. OpenStack Neutron APIs	323
Appendix C. REST API	327
Appendix D. Troubleshooting.	341
Log Information	341
Network Layers	341
Troubleshooting the Underlay Network	341
Troubleshooting the Management Network	342
Troubleshooting the Overlay Network.	342
Appendix E. Known Issues.	343
Appendix F. Upgrading IBM SDN VE Components	349
IBM SDN VE Controller	349
DOVE Connectivity Service (DCS)	349
DOVE Gateway (DGW).	350
Appendix G. Getting Help and Technical Assistance	351
Before You Call	351
Using the Documentation	351
Getting Help and Information on the World Wide Web	351
Software Service and Support	352
Hardware Service and Support	352
IBM Taiwan Product Service	352
Appendix H. Notices	353
Trademarks	353
Important Notes	354
Particulate Contamination	355
Documentation Format	356

Preface

This *User Guide* describes how to configure and use the IBM Software Defined Network for Virtual Environments (IBM SDN VE) version 1.2 to provide virtualization of the physical network using IBM Distributed Overlay Virtual Ethernet (DOVE) technology.

Who Should Use This Guide

This guide is intended for network installers and administrators engaged in configuring and maintaining a complex network. The administrator should be familiar with general Ethernet concepts and Layer 2 switching. They should also be familiar with the required VMware vCenter, vSphere, and ESX products and virtualization concepts.

What You'll Find in This Guide

This guide will help you plan, implement, and administer IBM SDN VE software. Where possible, each section provides feature overviews, usage examples, and configuration instructions. The following material is included:

Part 1: Getting Started

This material is intended to help those new to this product understand the basics of SDN VE installation and management. This part includes the following chapters:

- [Chapter 1, "IBM SDN VE Introduction,"](#) provides a conceptual overview of the SDN VE solution, and describes the prerequisites and general tasks for SDN VE installation.
- [Chapter 2, "Installing Unified Controller Modules,"](#) covers specific instructions for the installation and initial configuration of the SDN VE Controller software which provides the core intelligence of the SDN VE solution.
- [Chapter 3, "Installing DSA Modules,"](#) provides specific instructions for the installation and initial configuration of the Distributed Services Appliance (DSA) software which provides network connectivity to both virtual and physical network elements.
- [Chapter 4, "Installing the SDN VE 5000V Distributed vSwitch,"](#) provides specific instructions for the installation and initial configuration of the Distributed Switch 5000V software which provides virtual switching within a VMware virtual datacenter.
- [Chapter 5, "Virtual Network Configuration,"](#) provides specific instructions and examples for configuring elements of the virtual network.
- [Chapter 6, "Network Services,"](#) provides information on configuring network services such as logical groups, Layer 3 information, session monitoring, connectivity service configuration, and policy configuration.
- [Chapter 7, "Topology,"](#) provides a view of the topology and the interconnected switches and hosts in the logical groups and physical networks.
- [Chapter 8, "System Administration,"](#) provides information on managing system-related activities.

Part 2: Advanced Features

This section provides information on the advanced features in the SDN VE solution.

- [Chapter 9, “OpenStack,”](#) describes integration of OpenStack using IBM SDN VE Plugin.
- [Chapter 10, “Waypoint Connectivity Service,”](#) provides information on configuring middle boxes.
- [Chapter 11, “NIST,”](#) provides information on IBM SDN VE solution compliance with National Institute of Standards and Technology (NIST) Special Publication (SP) 800-131A specifications.
- [Chapter 12, “Public Key Infrastructure,”](#) provides information on the configuration of security and authentication.
- [Chapter 13, “Access Control Lists,”](#) describes how to use filters to permit or deny specific types of traffic, based on a variety of source, destination, and packet attributes.
- [Chapter 14, “Quality of Service,”](#) discusses Quality of Service (QoS) features, including Differentiated Services and IEEE 802.1p priority values.
- [Chapter 15, “sFlow,”](#) describes how to use the sFlow agent for sampling network traffic and providing continuous monitoring information to a central sFlow analyzer.
- [Chapter 16, “TCP Segmentation Offload,”](#) describes the use of TSO to reduce CPU overhead.
- [Chapter 17, “Virtual Router Redundancy Protocol,”](#) describes how VRRP can be used to configure high-availability.

Part 3: Command Reference

This section lists each command, together with the complete syntax and a functional description, from the Command-Line Interface (CLI).

- [Chapter 18, “Command Basics,”](#) provides an overview of the command syntax, including command modes and global commands.
- [Chapter 19, “Show Commands,”](#) provides a list of commands for collecting system configuration and statistics information.
- [Chapter 20, “Configuration Commands,”](#) provides a list of commands required to configure the virtual networks, and SDN VE components and features.
- [Chapter 21, “DSA Show Commands,”](#) provides an alphabetic list of Distributed Services Appliance (DSA) commands for collecting system configuration and statistics information.
- [Chapter 22, “DSA Configuration Commands,”](#) provides a list of DSA configuration commands.
- [Chapter 23, “Diagnostics Commands,”](#) provides commands to view diagnostic information.

Part 4: Appendices

- [Appendix A, “New and Updated Features,”](#) provides a summary of the updates in this release.
- [Appendix B, “OpenStack Neutron APIs”](#) provides a list of supported OpenStack Neutron APIs.
- [Appendix C, “REST API”](#) provides a list of supported REST APIs.
- [Appendix D, “Troubleshooting”](#) provides information on troubleshooting the SDN VE setup.
- [Appendix E, “Known Issues”](#) provides a list of known issues in the current release.
- [Appendix F, “Upgrading IBM SDN VE Components”](#) provides information on upgrading the SDN VE components.
- [Appendix G, “Getting Help and Technical Assistance,”](#) describes how to obtain product support.
- [Appendix H, “Notices”](#) includes the notices.

Typographic Conventions

The following table describes the typographic styles used in this book.

Table 1. *Typographic Conventions*

Typeface or Symbol	Meaning	Example
<i>ABC123</i>	This italicized body type shows book titles, special terms, or words to be emphasized.	Read your <i>User's Guide</i> thoroughly.
ABC123	This plain, fixed-width type is used for names of commands, files, and directories used within the body of the text. It also depicts on-screen computer output and prompts.	View the <code>readme.txt</code> file. <code>host#</code>
ABC123	This bold, fixed-width type appears in command examples. It depicts text that must be typed in exactly as shown.	<code>host# show config</code>
< >	Angled brackets appear in command examples as a parameter placeholder. Replace the indicated text with the appropriate real name or value when using the command. Do not type the brackets.	If the command syntax is: <code>ping <IPv4 address></code> You might enter: <code>ping 192.32.10.12</code>
[]	Square brackets depict optional elements within commands. These can be used or excluded as the situation demands. Do not type the brackets.	<code>host# ls [-a]</code>
{ A B }	Curled braces and vertical bars are used in command examples where there are multiple choices. Select only one of the listed options. Do not type the braces or bars.	If the command syntax is: <code>set {left right}</code> You might enter: <code>set left</code> Or: <code>set right</code>
AaBbCc123	This bold type depicts menus, buttons, and other controls that appear in Web browsers and other graphical interfaces.	Click the Save button.
A > B	This bold type with an angled right-bracket indicates nested menu items in a graphical interface.	Select File > Save .

How to Get Help

If you need help, service, or technical assistance, visit our web site at the following address:

<http://www.ibm.com/support>

The warranty card received with your product provides details for contacting a customer support representative. If you are unable to locate this information, please contact your reseller. Before you call, prepare the following information:

- Serial number of the switch unit
- Software release version number
- Brief description of the problem and the steps you have already taken
- Technical support dump information (`show tech dump support`)

Part 1: Getting Started

Chapter 1. IBM SDN VE Introduction

The IBM Software Defined Network for Virtual Environments (IBM SDN VE) version 1.2, is part of IBM's family of solutions for Software Defined Networking (SDN). IBM SDN VE components provide network virtualization within the IBM SDN platform while being agnostic to physical networks.

The IBM SDN VE consists of an architecture where applications, network services, and provisioning platforms can exploit the underlying network using a uniform API that is commonly known as the Northbound API. The IBM SDN VE architecture abstracts the underlying network and presents the network as a service or as an infrastructure.

SDN represents a major advance in enterprise communications. It creates a new network paradigm that separates network control logic from the underlying network hardware.

IBM SDN VE Solution Overview

The IBM SDN VE solution supplies a complete implementation framework for network virtualization. It supplies a core component of the SDN architecture, which is fully deployable for data center expansion. With SDN, instead of having to directly configure each connected device that makes up a network, administrators can dynamically establish multiple networks. They can allocate bandwidth and route data flows for optimized performance using high-level control programs.

The IBM SDN VE solution has a single point of control: the SDN Unified Controller. With the Unified Controller, resources can be abstracted and utilized in two ways:

- Overlay: unified network virtualization services based on IBM's Distributed Overlay Virtual Ethernet (DOVE) technology
- OpenFlow: logical groups (networks), based on OpenFlow technology

The IBM SDN VE, based on OpenFlow and IBM DOVE overlay technology, is a major part of the IBM's solution for SDN. By overlaying virtual networks onto physical networks, administrators can make existing infrastructure more adaptable to different workloads. The result is an agile, optimized, scalable network that is responsive to the needs of the business.

IBM SDN VE solution takes a host-based overlay approach, which achieves advanced network abstraction that enables application-level network services in large-scale multi-tenant environments. It provides a multi-hypervisor, server-centric

solution comprising multiple components that overlay virtual networks onto any physical network that provides IPv4 connectivity. The software is designed to support multi-vendor data center environments.

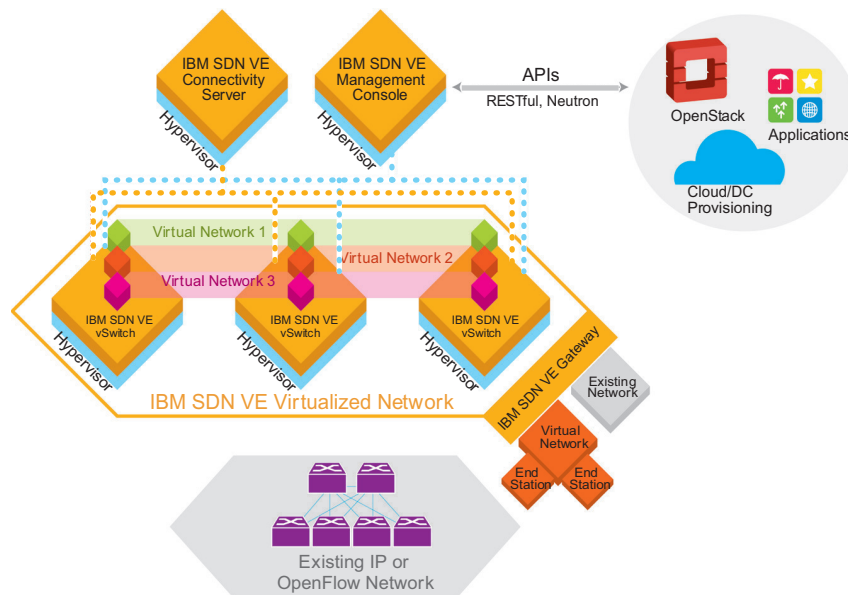


Figure 1. IBM SDN VE is a multi-hypervisor virtual network overlay that uses existing IPv4 infrastructure

Product Editions

The IBM SDN VE product suite is composed of the following editions:

- IBM SDN VE VMware Edition

This version is specifically targeted at running in conjunction with the VMware's vSphere hypervisor.

The IBM SDN VE VMware Edition requires an SDN VE Virtual Switch (an upgrade to the IBM Distributed Virtual Switch 5000V) to be resident in VMware. The IBM SDN VE VMware Edition is packaged for easy installation using VMware install and update tools.

- IBM SDN VE KVM Edition

This product version is targeted for the Kernel-based Virtual Machine hypervisor, or KVM in Linux.

The IBM SDN VE KVM Edition requires RHEL 6.5 server version of Linux which supports enhancements to VXLAN tunnels over Linux Bridge and an agent RPM, implementing SDN VE functionality to be installed.

- IBM SDN VE OpenFlow Edition

This product version is focused on providing IBM SDN VE Controller-integrated solution with support for versions 1.0 and 1.3 of the OpenFlow communications protocol.

OpenFlow is an emerging industry standard protocol that moves the network control plane into software running on an attached server. The IBM SDN VE OpenFlow Edition can be deployed in an environment with hosts connected to virtual and physical switches that have OpenFlow 1.0 or 1.3 protocol versions enabled.

Note: Although implementing the IBM SDN solution does not require changes to physical infrastructure, the hypervisor must be updated.

Note: The IBM SDN VE OpenFlow Edition can be used independently, or with the IBM SDN VE VMware or IBM SDN VE KVM Edition. In either case, you will have to purchase and install the OpenFlow Edition license. The IBM SDN VE OpenFlow solution can be implemented by deploying OpenFlow-enabled physical and virtual switches.

IBM SDN VE Solution Components

The IBM SDN VE solution is made up of four software components that work in combination to provide effective host-based network virtualization.

- **Unified Controller:** A controller provides the centralized point of control for configuring SDN VE that resides on a server as a virtual appliance. It allows administrators to manage individual networks and policies, and disseminates that virtual network and policy information to the connectivity service and gateways. The controller can be deployed in a highly available Active-Standby configuration.
- **Distributed Connectivity Service (DCS):** A connectivity service disseminates policies to the virtual switches participating in an SDN VE virtual network. The connectivity service software is deployed as a cluster of virtual appliances.
- **Distributed Gateways (DGW):** Enables SDN VE to establish interoperability with networks and servers that are external to the SDN VE environment.

Includes two gateways:

- Distributed VLAN Gateways: Enable VMs in an SDN VE tenant to connect with networks and servers that are external to the overlay network from a Layer 2 (VLAN) perspective.
- Distributed External Gateways:
 - Enable VMs in an SDN VE tenant to connect to non-SDN VE/DOVE external systems.
 - Enable VMs in an SDN VE tenant to connect to SDN VE VMs in another tenant through policy allocations.
 - Enable external systems to connect to VMs inside SDN VE/DOVE tenants.
- **5000V Host Module:** A Distributed Switch is software that resides in the hypervisor. It serves as the start and end point of each virtual network. The Distributed Switch provides Layer 2 and Layer 3 network virtualization over a UDP overlay, and implements the data path of the virtual network. The virtual switch also performs control plane functions to support virtual machine (VM) address auto discovery, VM migration and network policy configuration.

Unified Controller

The Unified Controller is a key component of the IBM SDN VE solution. It provides an abstracted view of the entire network and helps to manage the network and services.

The Unified Controller provides a rich set of application programming interfaces (APIs) that support multiple hypervisors across different hardware architectures. The IBM SDN VE product suite needs to be deployed once, after which

administrators can manage different hypervisors, network infrastructure, management policies, and vendor-dependent features when deploying added services to their respective environments.

The Unified Controller receives information from each SDN VE vSwitch, DCS, and Distributed Gateways. It is the central point to view operational and statistical information about the SDN VE solution components.

IBM SDN VE Additional Features

In addition to the Unified Controller, the IBM SDN VE solution includes the following features:

- **Logical Groups:** Provides the multi-tenancy service that create tenant-based logical groups. It provides northbound APIs that are compliant with OpenStack Neutron APIs v2.0.
- **Overlay Networks:** Supports overlay networks based on IBM DOVE technology for VMWare and KVM environments.
- **OpenStack Operation:** Supports OpenStack APIs.
- **Waypoint Connectivity Service:** Enables configuration and deployment of middle box service chains between logical groups.
- **Flow Replication and Redirection:** Enables logical SPAN service in an abstract manner over the fabric. This monitoring service provides a facility to replicate/redirect a subset of flows, based on session rules, between a source and a destination to a replication point.
- **Layer 3 Service:** Implements routing functionality in the OpenFlow network
- **DOVE Manager:** Establishes connection with the SDN VE Controller, forwarding service-level requests to a DOVE network, and handling notifications from the controller.
- **Static Flows Service:** Enables insertion of flows into a network.
- **Configuration and Monitoring:** Provides a Graphical User Interface (GUI) and Command-Line Interface (CLI) for configuration of the SDN VE components and features. System can be monitored with statistics and logging service, and topology visualization that provides logical and physical topology views.
- **Clustering:** Enables clustering of nodes that can be configured for high-availability.
- **Security:** Provides features, such as LDAP, RADIUS, and authentication, that help to secure the setup. IBM SDN VE also provides the option to be NIST SP 800-131A-compliant.
- **RBAC support:** Provides Role-based Access Control (RBAC) to restrict system access to authorized users.

Better and efficient forwarding of traffic through the fabric is enabled using the following services:

- **ARP Interposer:** Reduces Address Resolution Protocol (ARP) floods using controller-based proxy mechanisms.
- **Flow Merging and Conflict Resolution:** Enables multiple services to simultaneously run on the network.

IBM SDN VE Solution Benefits

The IBM SDN VE solution offers data center managers many ways to expand services and control costs. The solution helps to:

- Virtualize existing IPv4 networks with no change to the underlying physical network infrastructure
- Lower operating expenses by automating network provisioning and simplifying administration
- Expedite data center consolidation by allowing existing network addresses to be retained
- Enable large-scale multi-tenancy with independent management and optimization of multiple virtual networks
- Improve server resource utilization and return on investment by removing the network as a bottleneck to increase VM density
- Provide API-based programmatic access to virtual networks: data center provisioning platforms and network services can use virtual networks as a service or as an infrastructure

No Disruption to Existing IPv4 Networks

No CIO wants to replace a data center network. In most large-scale data centers, network administrators strive to wire the network one time then operate and maintain it without change. Changing the underlying physical infrastructure to support new business application requirements is hard to do and typically takes days or weeks to complete. This is a central problem data center managers must resolve. When compute and storage resources can be provisioned rapidly but network connectivity cannot, it can negatively impact business agility.

SDN VE helps data center managers increase business agility by enabling rapid provisioning of virtual network services without disrupting existing physical assets. The software does not require any change to existing networks to operate—a valuable attribute that simplifies adoption. The only requirement to implement SDN VE is a simple one. The physical network infrastructure on which the software is overlaid must be capable of providing IPv4 address-based connectivity. Every typical enterprise data center network supports this capability.

SDN VE efficiently overlays virtual networks onto existing networks, thus decoupling application connectivity from the physical network infrastructure. This enables a “wire once” physical network that can support multiple SDN VE virtual networks flexibly managed and controlled through highly available clusters.

IBM SDN architecture separates the control plane from the data plane, a central tenet of SDN. SDN VE operates by adding a distinct header to packets sent by VMs. Each SDN VE data transfer is just an ordinary IPv4 packet sent to the existing switches in the data center network and the switches can use existing IPv4

forwarding routes and tables. Devices continue to operate at line rates. The IBM SDN VE solution builds on the network that is already in place, and provides the flexibility to create and manage virtual networks on demand.

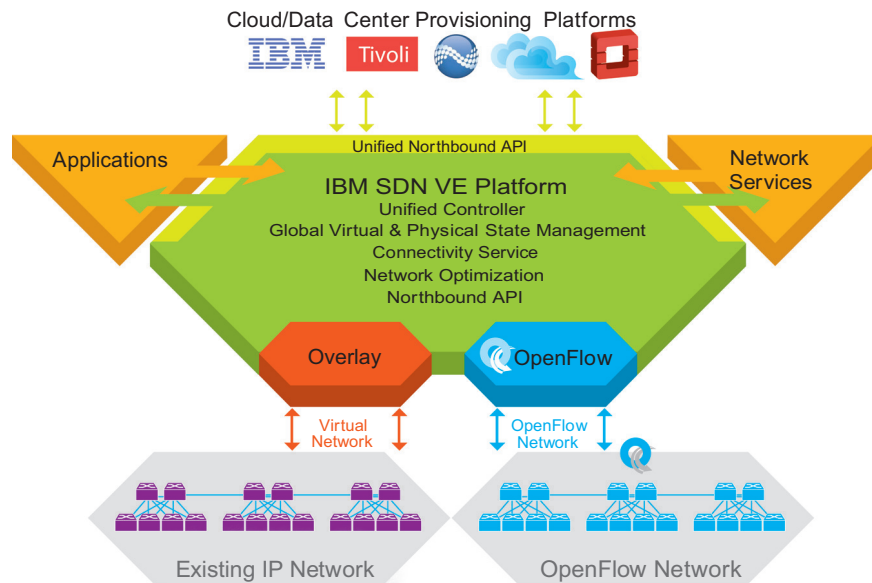


Figure 2. IBM SDN VE abstracts the underlying network and presents it to applications as either a service or as an infrastructure

Extending the VM Strategies into the Network

SDN VE is a logical extension of the virtualization trend that has become the dominant feature in the data center. The software extends the efficiency and productivity advantages achieved with server virtualization to the process of network provisioning and management. These advantages allow data centers to be more:

- Efficient, because SDN VE improves resource use. It allows secure, dedicated virtual networks to be created quickly and easily, without requiring changes to the underlying physical infrastructure.
- Agile, because SDN VE cuts network provisioning time from days to minutes. With SDN VE, you can establish secure virtual networks as easily as starting up VMs.
- Scalable, because SDN VE offers data center managers the scalability needed for current and future growth. Up to 16 million networks can be specified in the architecture. This release of IBM SDN VE supports up to 16,000 virtual networks.

High Availability

Enterprise data centers maintain uncompromising standards for high availability, which reflects the value that data center operations contribute to the enterprise. In many cases, the data center is one of the most valuable components in the business because the enterprise cannot function if the data center is down. SDN VE supports enterprise needs for high availability with customizable, redundant component design.

Connectivity Service

In SDN VE, virtual networks are collected into administrative constructs called tenants. A connectivity service disseminates VM addresses to the virtual switches participating in an SDN VE virtual network. The connectivity service software is deployed as a cluster of virtual appliances. Two or more active SDN VE Connectivity Servers control each virtual network within a tenant. The number of SDN VE Connectivity Servers that can be assigned to individual tenant is user-configurable. This ensures that the user can select the level of high availability needed for a given virtual network. This redundant design allows the state of each SDN VE Connectivity Server to be replicated in at least one other instance of the SDN VE Connectivity Server at all times.

Unified Controller

The SDN VE Controller provides high availability in Active and Standby modes. One instance operates in Active mode, and the other functions in Standby mode. If an Active SDN VE Controller experiences a failure or outage, automatic failover to the Standby SDN VE Controller occurs.

Enhanced Multi-tenancy for Cloud Providers

The gains in adopting SDN VE are far greater than employing VLANs. With SDN VE, you can create secure, scalable multi-tenant networks with individual network control. Each virtual network created with SDN VE can be managed individually using the application programming interface (API) the software provides. In addition, you get greater scalability with SDN VE: A traditional network is physically limited to 4096 VLANs, and requires configuration of end-to-end VLANs on some or all physical devices in the network. With SDN VE, the maximum number of VLANs that can be supported increases from a physical limit of 4096 networks to an architectural limit of 16,000,000. This release of IBM SDN VE VMware Edition supports 16,000 virtual networks. Cloud providers that need to support multiple customers with dedicated, reliable, secure and scalable networks, can deploy SDN VE to help supply these services with increased cost effectiveness and efficiency.

Datacenter Consolidation

Datacenter consolidation is a common practice among large enterprises today because of the increased economy and efficiency that can be gained. One difficulty of consolidation is combining IPv4 addresses. Redesigning complete network schemas is an exceptionally complex and time-consuming task. SDN VE resolves this problem by reusing existing IPv4 addresses. Each logical group can have overlapping IP addresses. Only the MAC addresses of all the VMs need to be unique.

Maximizing Servers

VMs require real network connections. However, since it is much easier to create VMs than it is to network them, your network resources can be exhausted before you can use your servers to the fullest extent. Maximizing server use is a principal reason to implement SDN VE. With the software in place, VM density can be increased to the limits of memory, and processor cycles and server virtualization can continue without concern for VM network bottlenecks. With SDN VE, you can establish a "wire-once" data center network environment with expansion capacity for future growth and increased virtualization.

Optimizing Provisioning with Programmable APIs

The IBM SDN VE solution provides programmatic access to virtual network functions using RESTful APIs, which can provide web services to any client program able to transmit messages using the HTTP or HTTPS protocols. SDN VE also supports the OpenStack Neutron API, which is a network abstraction that allows OpenStack to use the underlying network as the infrastructure without requiring it to have knowledge of the underlying resources.

IBM SDN VE Solution Elements

The IBM SDN VE solution requires the following components:

SDN VE Elements (all Editions)

- Unified Controller

This IBM software resides on two VMware VMs on different hosts within the virtual datacenter. Together, they provide the resilient core intelligence for DOVE, unifying the operation of various VM-based service appliance modules that form the fabric of the distributed virtual network.

- Distributed Services Appliance (DSA)

This IBM software resides in multiple VMware VMs. Each has the capacity to become a DCS or a DGW as described below.

- Distributed Connectivity Service (DCS)

These IBM software modules collect and process network information pertaining to nearby VMs, gateways and virtual switches in the virtual datacenter. Tenant information is synchronized among partner modules within the distributed virtual network.

- Distributed Gateway (DGW)

These IBM software modules can serve as a gateway to join the virtual network to an external, non-virtual network associated either with a specific port in the physical network or with legacy VLAN broadcast domains.

VMware Edition Elements

- VMware vCenter

This VMware product resides on a server within the datacenter. It provides a centralized tool for installing, managing and synchronizing hypervisors, virtual machines (VMs), and virtual distributed switches (vDS) on host servers throughout the datacenter.

- VMware vSphere Client or vSphere Web Client

This VMware vSphere Client resides on administrative client devices. It provides the server administrator or network administrator with rich, remote access to vCenter management tools. The vSphere Web Client provides similar access via your web-browser interface.

- VMware ESX 5.0 or 5.1 or 5.5

These VMware hypervisor products reside on individual host servers within the datacenter. They provide the software infrastructure for installing, running, and managing VM and vDS elements on the hosts.

- SDN VE 5000V

The 5000V is a versatile vDS solution. Though it can be used independently to provide general virtual switching within a VMware virtual datacenter (outside of the IBM SDN VE solution), it is a required element within SDN VE solution:

- 5000V Host Modules

This IBM software resides in participating VMware ESX hypervisors on host servers within the virtual datacenter. It implements a vDS portset as defined in the VMware vDS API and acts a virtual network switch for the given host server. At its core, it forwards frames based on destination MAC addresses, controlling Layer 2 access to and from the associated VMs. It also provides advanced switching features such as VLANs, IGMP snooping, etc. In the IBM SDN VE solution, the 5000V vDS host modules act as Tunnel End-Points (TEPs).

- 5000V Controller

This IBM software resides in a VM within the datacenter. It works in conjunction with SDN VE and VMware modules to unify the 5000V host modules associated with a specific vDS into an aggregate superswitch.

Prerequisites

The following must be provided prior to SDN VE installation:

- VMware vCenter Server must be installed and operational in your network (see the documentation provided with your vCenter product).
- All host servers which take part of the IBM SDN VE solution must be installed and operational, and include the following:
 - There should be at least three hosts for vMotion: at least two control nodes and at least one compute node.
 - Each host must have a minimum of one 1G or 10G physical NIC.
 - Each host must have IPv4 Layer 2/Layer 3 network connectivity to the vCenter and all host servers which will participate in their virtual network domain. IPv6 is not presently supported.
- In addition to the general host requirements:
 - Each host server that includes a SDN VE Controller, DSA, or 5000V vDS host module must have ESX 5.0 or 5.1 installed and operational.
 - The host server that includes the 5000V Controller, it is highly recommended that VMware High Availability and/or VMware Fault Tolerance features be configured to protect the virtual switch against downtime or data loss.
 - Each host server that includes a 5000V vDS host module must also have a valid VMware Enterprise Plus license installed.

- VMs for SDN VE Controller, DCS, and DGW modules must include the following:
 - For SDN VE Controller: Two VMs on different ESX hosts are required.
 - For DCS: Two VMs on different ESX hosts are required (three are recommended).
 - For DGW: Two VMs on different ESX hosts are required.
 - For the 5000V Controller, one VM is required.
 - For the 5000V vDS host module, one VM is required for each host that will include a vDS portset.
 - Each VM used as SDN VE entity must have a minimum allocation of 8 GB of memory.

The following SDN VE software files are required:

- Open Virtual Appliance (OVA) files for—
 - Unified Controller
 - Distributed Services Appliance (DSA)
 - DOVE Virtual Switch
 - 5000V Host Module
- VIB offline bundle file for the SDN VE 5000V Distributed vSwitch for VMware vSphere. This file includes the vSphere Installation Bundle (VIB).

Installation Summary

The following tasks summarize the SDN VE installation process and are covered in detail in the installation chapters:

Installing Unified Controller Modules

- Using VMware vSphere to deploy the SDN VE OVA file to VMs on two hosts.
- Initial SDN VE setup, including:
 - Starting the modules
 - Logging in to the CLI
 - Setting each module's IPv4 parameters
 - Establishing high-availability for system resilience

Installing DSA Modules

- Using VMware vSphere to deploy the DSA OVA file on at least five VMs.
- Initial DSA setup, including:
 - Starting the modules
 - Logging in to the CLI
 - Setting each module's IPv4 parameters
 - Attaching the modules to the SDN VE cluster
 - Specifying a connectivity or gateway role for each module
 - Setting the DOVE tunnel IP for the DSA module in gateway role

Installing the SDN VE 5000V Distributed vSwitch for VMware vSphere

- Using VMware vSphere to deploy the DS 5000V Controller OVA file on a VM.
- Initial 5000V setup, including:
 - Starting the module
 - Logging in to the CLI
 - Setting the module's IPv4 parameters
 - Creating a global vDS instance in the vCenter
 - Attaching the module to the SDN VE cluster

Configuration Summary

The following tasks summarize the SDN VE configuration process and are covered in detail in the network configuration chapter:

Configure the overlay network

- Create tenants
 - Create connectivity groups
 - Define network address space
 - Define policies
 - Export connectivity group configuration to the virtual switch (required only if OpenStack is not used)

Externalize the overlay networks

- Configure VLAN Gateways
- Configure External Gateways
- Configure Gateway Interfaces (required only if OpenStack is not used)

Configure Virtual Switch Host Module

- Install virtual switch Host Modules
- Configure the Underlay (Physical) Networks
- Attach ESXi Hosts to the virtual switch
- Configure Tunnel End-Points

Attach End Systems

Chapter 2. Installing Unified Controller Modules

Note: The instructions in this chapter are for new installations only. When upgrading existing DOVE components, please refer to Appendix C.

The Unified Controller provides the core intelligence that unifies the operation of the individual appliance modules installed on the participating host servers. Unified Controller modules must be installed and initialized on two different hosts for high-availability (HA) resilience as covered in this chapter.

Deploying the Unified Controller Software

The Unified Controller is installed as a virtual appliance. The VM image is provided in `raw` and `qcow2` formats, and in `vmdk` format. Use the image that suits your setup environment.

It should be deployed in the following minimal VM configuration or better configuration.

- 4 core CPU
- 8 GB RAM allocation for VM
- Cluster configuration of two-four nodes
- 16 GB HD

Install Unified Controller on Host

Note: **Two Unified Controller modules on different hosts are required.** Perform the steps in this section once for a *primary* Unified Controller, and again for a *secondary* Unified Controller on another host.

Note: You must enter the same license information on both the primary and secondary controller modules.

KVM Environment

Follow these steps to install the Unified Controller module:

1. Download the “SDN VE Unified Controller” image.
Image file name example: `SDNVE_UnifiedController_S4_041.kvm`
2. Untar the image: `#tar xvf SDNVE_UnifiedController_S4_041.kvm`
You will see two files with names similar to the following:
`ibmSDN-disk1.qcow2`
`ibmSDN-disk2.qcow2`
3. Save the two `.qcow2` files into the path: `/var/lib/libvirt/images:`
`#mv ibmSDN-disk?.qcow2 /var/lib/libvirt/images/`
4. Change path to `/etc/libvirt/qemu.`
5. Create `SDN-template.xml` with following contents:

```
# vi SDN-template.xml
```

```

<domain type='kvm' id='10'>
  <name>SDN_VE</name>
  <memory>8388608</memory>
  <currentMemory>8388608</currentMemory>
  <vcpu>2</vcpu>
  <os>
    <type arch='x86_64' machine='rhel6.5.0'>hvm</type>
    <boot dev='hd' />
  </os>
  <features>
    <acpi />
    <apic />
    <pae />
  </features>
  <clock offset='utc' />
  <on_poweroff>destroy</on_poweroff>
  <on_reboot>restart</on_reboot>
  <on_crash>restart</on_crash>
  <devices>
    <emulator>/usr/libexec/qemu-kvm</emulator>
    <disk type='file' device='disk'>
      <driver name='qemu' type='qcow2' />
      <source file='/var/lib/libvirt/images/ibmSDN-disk1.qcow2' />
      <target dev='vda' bus='virtio' />
      <alias name='virtio-disk0' />
      <address type='pci' domain='0x0000' bus='0x00' slot='0x07'
function='0x0' />
    </disk>
    <disk type='file' device='disk'>
      <driver name='qemu' type='qcow2' />
      <source file='/var/lib/libvirt/images/ibmSDN-disk2.qcow2' />
      <target dev='vdb' bus='virtio' />
      <alias name='virtio-disk1' />
      <address type='pci' domain='0x0000' bus='0x00' slot='0x08'
function='0x0' />
    </disk>
    <controller type='ide' index='0'>
      <alias name='ide0' />
      <address type='pci' domain='0x0000' bus='0x00' slot='0x01'
function='0x1' />
    </controller>
    <controller type='virtio-serial' index='0'>
      <alias name='virtio-serial0' />
      <address type='pci' domain='0x0000' bus='0x00' slot='0x05'
function='0x0' />
    </controller>
    <interface type='direct'>
      <source dev='eth2' mode='bridge' />
      <model type='virtio' />
      <address type='pci' domain='0x0000' bus='0x00' slot='0x03'

```



```

function='0x0' />
</interface>
<serial type='pty'>
  <source path='/dev/pts/2' />
  <target port='0' />
  <alias name='serial0' />
</serial>
<console type='pty' tty='/dev/pts/2'>
  <source path='/dev/pts/2' />
  <target type='serial' port='0' />
  <alias name='serial0' />
</console>
<channel type='spicevmc'>
  <target type='virtio' name='com.redhat.spice.0' />
  <alias name='channel0' />
  <address type='virtio-serial' controller='0' bus='0' port='1' />
</channel>
<input type='mouse' bus='ps2' />
<graphics type='vnc' port='5901' autoport='yes' />
<sound model='ich6'>
  <alias name='sound0' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x04'
function='0x0' />
</sound>
<video>
  <model type='qxl' vram='65536' heads='1' />
  <alias name='video0' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x02'
function='0x0' />
</video>
<memballoon model='virtio'>
  <alias name='balloon0' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x06'
function='0x0' />
</memballoon>
</devices>
</domain>

```

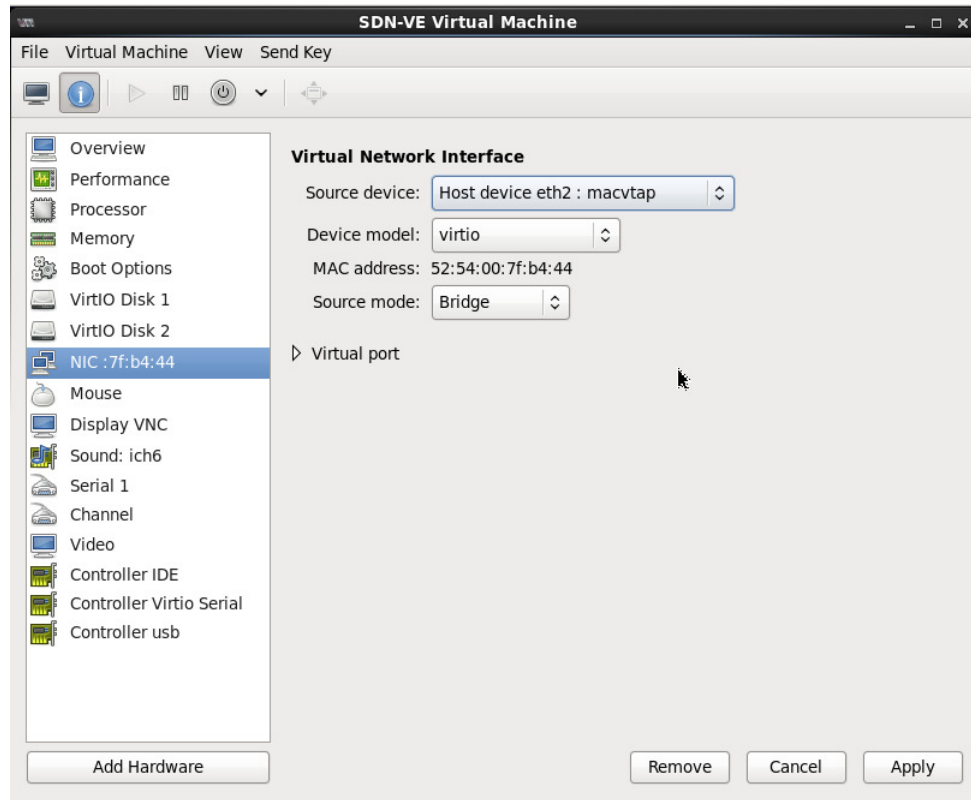
6. Run command:

```
[root@rhel65-28 qemu]# virsh define SDN-template.xml
```

If successful, you will see the following message:

```
Domain SDN VE defined from SDN-template.xml
```

7. Use virt-manager to edit the VM settings and map the NIC to the network intended as the controller's management network.



VMware Environment

Install Using OVF Tool

Use the following command to download and install the OVA file:

```
$/opt/vmware/ovftool/ovftool --name=<NAME> [--powerOn]
--datastore=<"NAME"> --network="<Network>" dmc.ova vi://<vCenter
IP>/<DC>/Host?ip=<Host IP>
```

Replace the variables in the command with appropriate values as follows

Table 2. Command Parameters

Option	Description
Name	Name of the VM. For example: SDN-VE-1.
powerOn	Specify this argument if want the controller to power on when the command is executed.
Datastore	Specify name of datastore.
Network	Select the VM network.

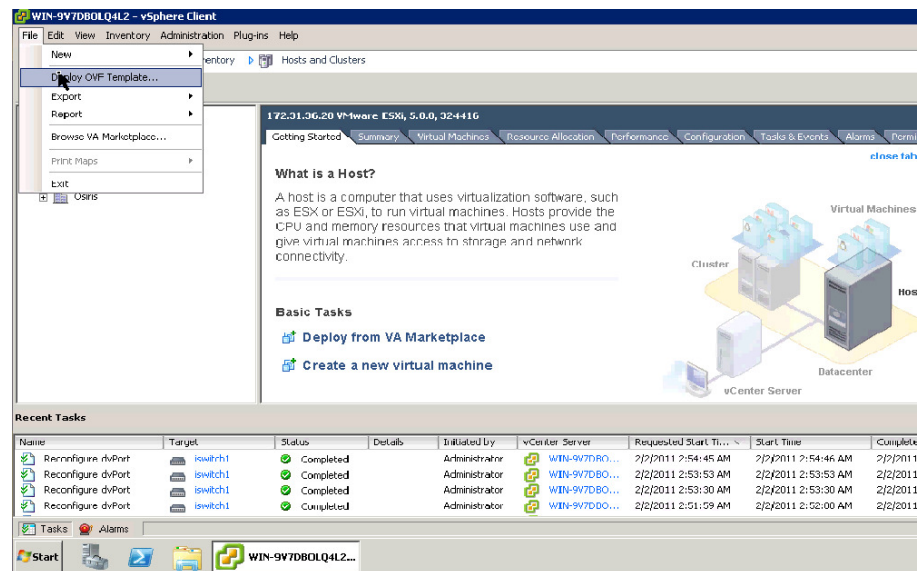
You will be prompted to enter the target host login credentials.

If you want to use vSphere Client to power on the controller, log in to the vSphere Client and follow [Step 13](#) onwards on [page 39](#).

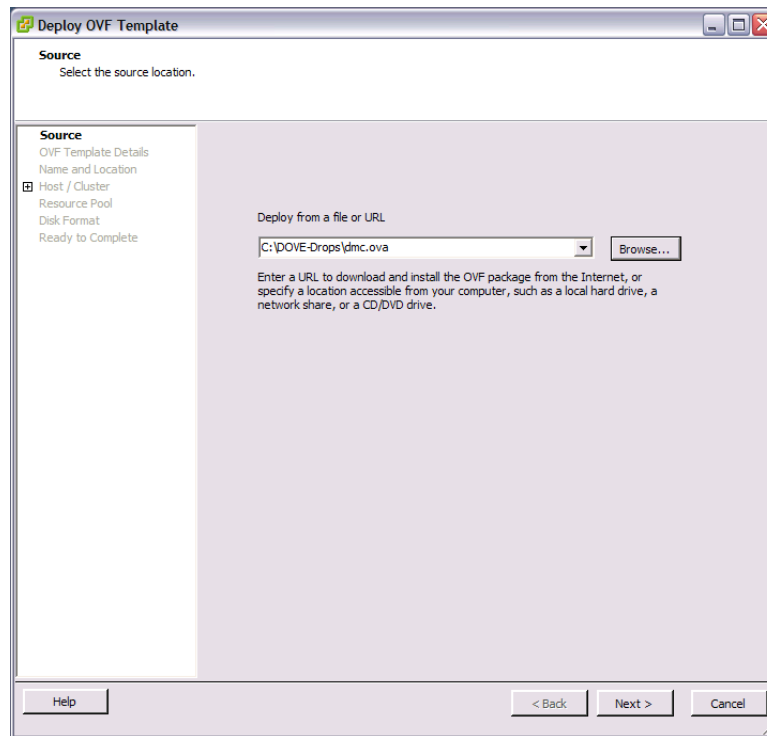
Install Using VMware vSphere Client

Follow these steps to install the controller module:

1. Download the controller image.
Image file name example: SDNVE_UnifiedController_S4_041.ova
2. Place the OVA file on a system that has access to the VMware vSphere Client (such as an administrative laptop).
3. Select an ESX host on which to deploy the controller.
The controller host merely provides an environment in which the SDN VE system will run. It is not required to participate as a vDS host and may be a different class of device than those where the vDS host modules will be installed. The primary requirement is for the controller host to have Layer 3 connectivity to the designated vCenter and participating DSA modules.
4. Launch the VMware vSphere Client and connect either to the vCenter that manages the host where the controller will be deployed or directly to the ESX host.
5. From the vSphere Client, select the target ESX host and choose **File > Deploy OVF Template** as shown below:

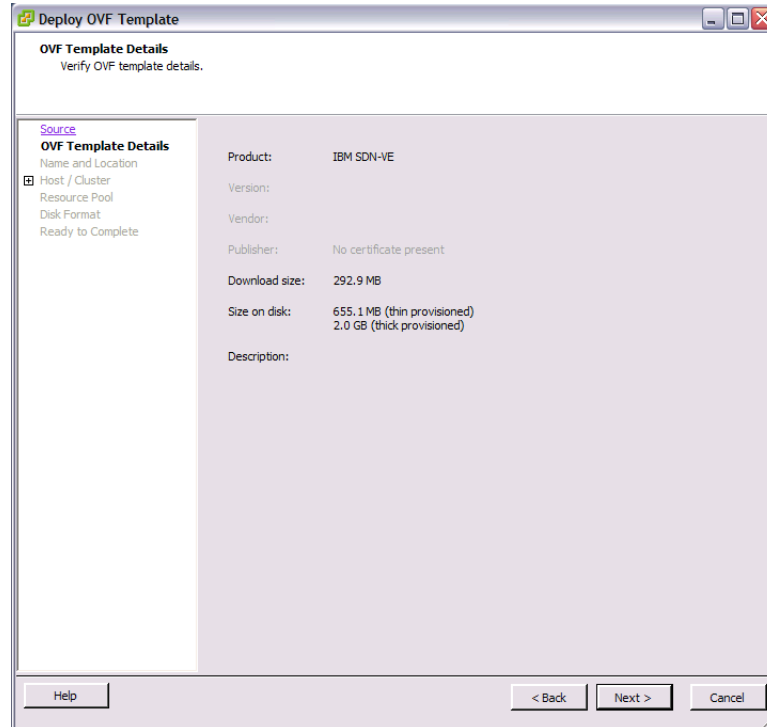


6. Select the location where the OVA file is stored and click **Next**.




The 'Deploy OVF Template' dialog box is shown in the 'Source' step. The title bar reads 'Deploy OVF Template'. The main heading is 'Source' with the instruction 'Select the source location.' On the left, a tree view shows 'Source' selected, with sub-items: 'OVF Template Details', 'Name and Location', 'Host / Cluster', 'Resource Pool', 'Disk Format', and 'Ready to Complete'. The main area is titled 'Deploy from a file or URL' and contains a text box with 'C:\DOVE-Drops\dmc.ova' and a 'Browse...' button. Below this, a note states: 'Enter a URL to download and install the OVF package from the Internet, or specify a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.' At the bottom are 'Help', '< Back', 'Next >', and 'Cancel' buttons.

7. Verify the OVA details and click **Next**.




The 'Deploy OVF Template' dialog box is shown in the 'OVF Template Details' step. The title bar reads 'Deploy OVF Template'. The main heading is 'OVF Template Details' with the instruction 'Verify OVF template details.' On the left, the tree view shows 'OVF Template Details' selected, with sub-items: 'Source', 'Name and Location', 'Host / Cluster', 'Resource Pool', 'Disk Format', and 'Ready to Complete'. The main area displays the following details: Product: IBM SDN-VE; Version: (blank); Vendor: (blank); Publisher: No certificate present; Download size: 292.9 MB; Size on disk: 655.1 MB (thin provisioned) and 2.0 GB (thick provisioned); Description: (blank). At the bottom are 'Help', '< Back', 'Next >', and 'Cancel' buttons.

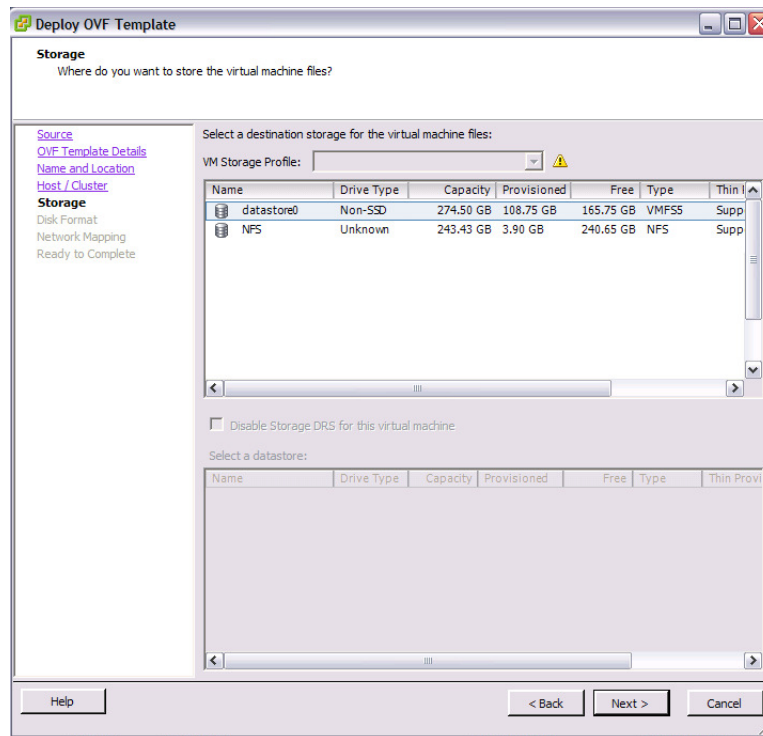
8. Provide a name for the controller module and click **Next**.



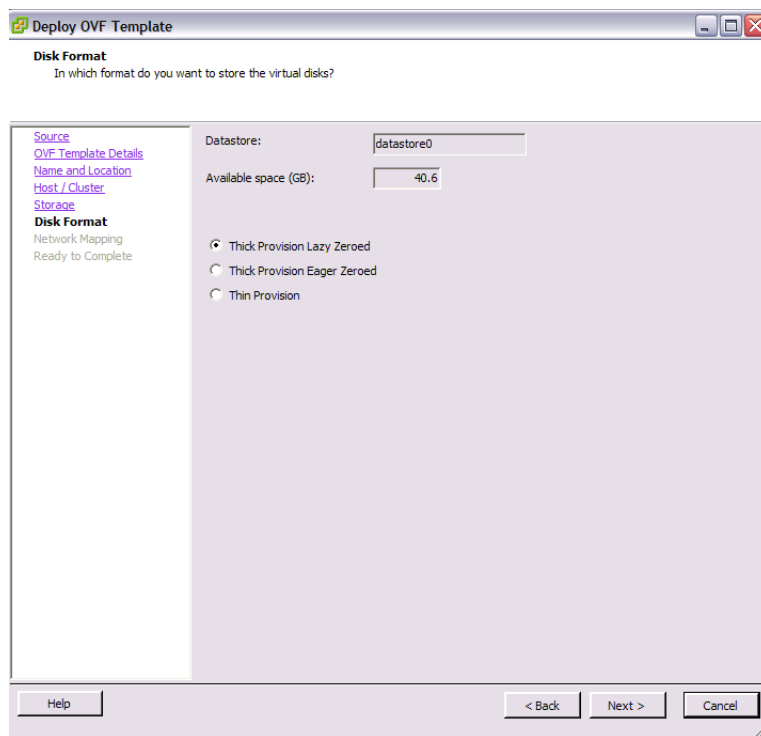
9. Specify the host or cluster on which to deploy the controller and click **Next**.



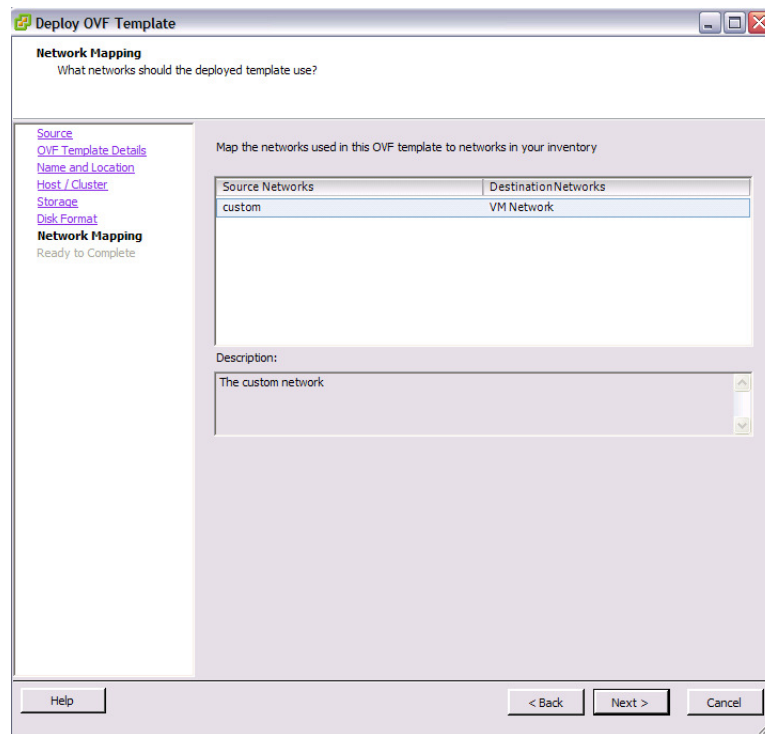
10. Specify a location on the VM where controller files should be stored, and click **Next**.



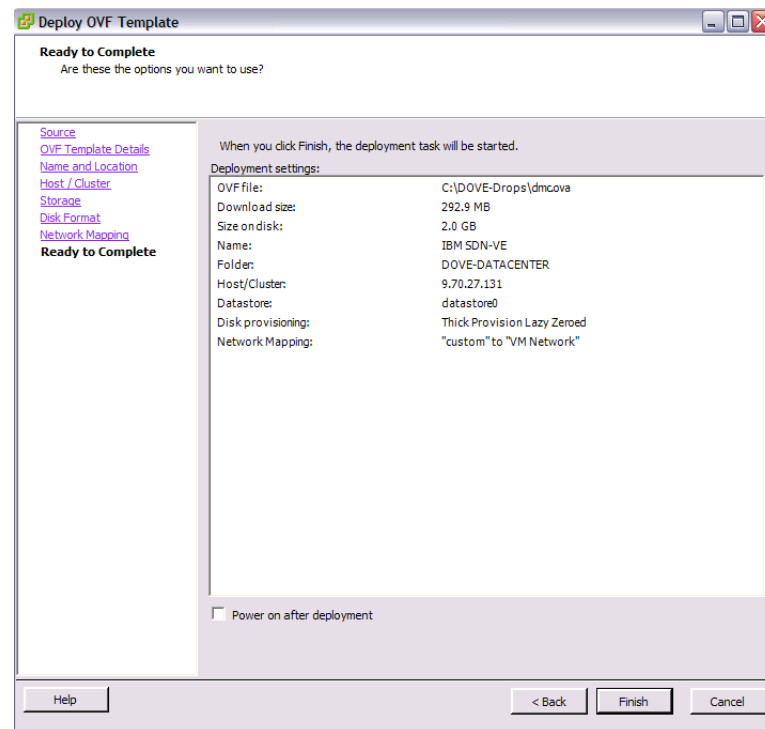
11. Select a disk format and click Next. The recommended format is Thick Provisioned Lazy Zeroed.




12. Map the network for controller use and click **Next**.



13. Verify the specified options, select the “Power on after deployment” option, and click **Next**.



This will initiate the controller module VM deployment:



The controller VM will power on when deployment is complete, and the IBM SDN VE management console will appear.

Initial Unified Controller Setup

After installing primary and secondary Unified Controller module software, each Unified Controller must be manually configured by entering commands into the built-in Command-Line Interface (CLI).

Perform the following initial Unified Controller setup for both the primary and secondary Unified Controller modules.

Start the Unified Controller Module

KVM Environment

Initially, the CLI can be accessed only through the virt-manager. Later, if desired, the CLI can be accessed via remote SSH connection or configuration can be performed via the Graphical User Interface (GUI).

Note: The Unified Controller module can be started using virt-manager, the `virsh start` command, or by accessing the module using a remote connection. The steps in this section are for using the virt-manager.

In virt-manager:

1. Right-click on the Unified Controller and select Run. The Unified Controller is powered on.
2. Select the **Open** icon.
3. Select the **Console** icon to open the Unified Controller CLI.


VMware Environment

Initially, the CLI can be accessed only through each controller VM console on the vSphere Client. Later, if desired, the CLI can be accessed via remote SSH connection or configuration can be performed via the Graphical User Interface (GUI).

When following the provided installation instructions, the controller module automatically starts when the VM is powered on.

However, to manually access the console under other conditions, use the following procedure:

1. Log-in to the VMware vCenter via your vSphere Client.
2. Right-click on the controller VM and select the option to “Open Console.” Alternately, you can click on the Console icon.



The VM console for the controller will appear.

Set the Language

When the Unified Controller CLI opens, you will be prompted to set the language:

```
SET_ADDRESS> language symbol en_US
```

Set the Unified Controller IPv4 Address (Optional)

Note: This section is not required if you are using DHCP.

Each Unified Controller must have IPv4 connectivity to the hosts that will participate in the SDN VE system.

By default, the Unified Controller is enabled for dynamic IPv4 addressing using DHCP. If there is a DHCP server available in your network, the Unified Controller will automatically acquire IPv4 address.

However, if DHCP is not available in your network or if you wish to override DHCP and configure static IPv4 addresses for the Unified Controller, enter the following commands, depending on whether you prefer IPv4 address/netmask or CIDR notation.

If using static IP configuration, set the Unified Controller IPv4 address using the following command

Using IPv4 Address and Netmask:

Using IPv4 Address and Netmask

The following steps use IPv4 dotted-decimal (*a . b . c . d*) notation.

1. Set the Unified Controller address:

```
SET_ADDRESS> ipmgmt set ip addr <Unified Controller IPv4 address> mask  
<netmask>
```

2. Optional. Set a gateway (router/next-hop) address:

```
SET_ADDRESS> nexthop set ip addr <gateway address> mask <netmask>
```

(OR)

Using CIDR Notation

The following steps use CIDR dotted-decimal (*a.b.c.d/e*) notation.

1. Set the Unified Controller address:

```
SET_ADDRESS> ipmgmt set cidr <Unified Controller address>
```

2. Optional. Set a gateway (router/next-hop) address:

```
SET_ADDRESS> nexthop set cidr <gateway address>
```

Using DHCP

DHCP is used by default. However, if you have configured static IPv4 addresses and prefer to return to DHCP operation, enter the following command:

```
SET_ADDRESS> ipmgmt set dhcp
```

Note: Switching to DHCP will clear the static IPv4 addresses for the Unified Controller and its gateway, DNS, and high-availability configuration.

Verifying Addresses

You can verify Unified Controller IPv4 address and gateway configuration using the `show ipmgmt` command.

You can verify DNS settings using the `show dns` command.

Enter License Information

Note: You must enter the same license information on both the primary and secondary controller modules.

After specifying IP address and language, exit the initial configuration CLI.

In a browser, specify the following URL to log in to the controller.

`https://<Unified Controller IPv4 address>:8443`

Log in using the default credentials:

Default user name: **admin**

Default password: **admin**

You will be prompted to enter the 64 character license key and to accept the license text. After license key acceptance, the appliance will complete its startup.

Note: You may add as many licenses as you need to. Or, you may add additional licenses at any point in time using the controller GUI.

Establish SDN VE Controller HA

When HA is established, all configuration has to be performed via the active controller. The configuration is automatically synchronized with the standby controller, which will take over as the active controller if the initial active controller fails.

On the GUI:

1. Access the SDN VE HA page:
Administration > System Tools > SDN-VE HA.

Node	Role	Local Host	Status	Sync Status
9.121.62.116	Standby	true	online	completed
9.121.62.118	Active	false	offline	na

2. Select **Add to Cluster**.
3. Enter the Cluster Name and IP address(es) of the cluster node(s) – comma separated in case of multiple IP addresses.

Add to Cluster

Cluster Name:

Cluster Nodes:

OK Cancel

If the cluster is successfully added, a “Cluster configuration completed” message is displayed on the top left corner.

Note:

- All the controllers that are part of a cluster must have the same configuration and license information.
- When the operation “Add to Cluster” is performed for the very first time, you must restart the node once using the command **Administration > System Tools > System Commands > reboot VM**. Subsequent “Add to Cluster” operations do not require a restart.
- Configure all the controller nodes as Primary.

You can view the cluster information on the **Administration > System Tools > SDN-VE HA** page.

Configuration changes response will be shown here...

System Tools

- System Commr
- SDN-VE HA**
- Log
- License
- LDAP
- NIST
- Openflow
- RADIUS
- Upgrade
- Back-up / Rest
- SDN-VE DOVE
 - Authenticat
 - HA
 - IP Mgmt
 - IPSEC
 - LG Global
 - Controller PKI

SDN-VE HA Settings

DMC : Members

Node	Role	Local Host	Status	Sync Status
9.121.62.116	Standby	true	online	completed
9.121.62.118	Active	false	offline	na

Rejoin Cluster Add to Cluster Disconnect from Cluster

Note:

- After the SDN VE HA configuration is complete, one of the two cluster nodes is set as Active and the other as Standby.
- The status (online / offline) indicates the cluster view status from the specific controller’s perspective and not the ping / reachability status.
- Only offline primary nodes shall appear and shall not include offline non-primary nodes.

Log in to the Unified Controller

Note: You can use the GUI to specify the settings required to complete the installation. The use of CLI is documented in this *User Guide*.

Access the CLI using a remote SSH connection, or via virt-manager.

Once you are connected to the Unified Controller, you are prompted to enter a login name and password. The default log-in user name and password are as follows:

Default user name: **admin**

Default password: **admin**

It is recommended that you change all default system password after initial configuration and as regularly as required under your network security policies.

Enter Configuration Mode

The Unified Controller uses a rich CLI command set with multiple command modes. For an overview of CLI modes and features, see [“Command Basics” on page 187](#). The remainder of this chapter will display all commands necessary for initial configuration, but only those command relevant to the specific configuration tasks will be called into example.

After logging in, perform the following commands to enter the CLI Global Configuration mode:

```
SDN-VE@SDN-VE-Controller# configure terminal
SDN-VE@SDN-VE-Controller(config)#
```

Establish Unified Controller High-Availability

As noted in the preceding installation process, two Unified Controller modules on different hosts are required for high-availability (HA). HA provides resilience in the event that the active Unified Controller fails. When HA is established, all configuration has to be performed via the primary Unified Controller. The configuration is automatically synchronized with the secondary Unified Controller, which will take over as the active Unified Controller if the primary fails.

Configure SDN VE DOVE HA

Note: Before proceeding with this section, ensure you have completed the steps in section [“Establish SDN VE Controller HA” on page 43](#).

On the primary Unified Controller module, configure an external IPv4 address that will be used for master access to the Unified Controller HA cluster as a whole. The address can be configured by IPv4 address and netmask or CIDR designation using one of the following Global Configuration mode commands:

1. Access the Unified DOVE configuration mode:

```
SDN-VE@SDN-VE-Controller(config)# sdnve-dove terminal
```

2. Configure external IP:

```
SDN-VE@SDN-VE-Controller (config-sdnve-dove)# external-ip ip <IPv4 address>
mask <subnet mask>
```

Note: No reboot is required when you first configure the external IP address. However, if you change the external IP address anytime later, you must reboot both the controllers using the command `reboot VM`.

3. Configure peers:

Note: You must configure the Active controller IP address (See [“Establish SDN VE Controller HA” on page 43](#)) as the primary SDN VE Controller.

```
SDN-VE@SDN-VE-Controller (config-sdnve-dove)# peers primary <IPv4 address of
primary SDN-VE Controller> mask <subnet mask of primary SDN-VE Controller> secondary <IPv4
address of secondary SDN-VE Controller> mask <subnet mask of secondary SDN-VE Controller>
```

4. Synchronize HA:

```
SDN-VE@SDN-VE-Controller (config-sdnve-dove)# ha-synchronization start
```

5. Verify HA synchronization using the following command:

```
SDN-VE@SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove
sync-status
```

6. Start HA:

```
SDN-VE@SDN-VE-Controller (config-sdnve-dove)# ha start
```

7. Verify the HA setting with the following command:

```
SDN-VE-Controller(config)# show sdnve-dove ha

ha          running
external   ip 9.121.62.240 mask 255.255.254.0
primary    ip 9.121.62.118 mask 255.255.254.0
secondary  ip 9.121.62.116 mask 255.255.254.0
```

The Unified Controller setup is now ready.

The Graphical User Interface

Most of the common configuration, management, and operation features of the SDN VE can be accessed via the Graphical User Interface (GUI) using a standard Web browser.

The GUI supports HTTPS on default port 8443 and is available once initial configuration of Controller HA is complete. To access the GUI, enter the following URL into your browser:

```
https://<Controller HA external IPv4 address>:8443
```

Note: Be sure to use the HA external IPv4 address for the Controller cluster, and not the individual primary or secondary controller IPv4 address. This helps ensure connection in case the primary controller fails.

Next Steps

Once high-availability is operating on the controller cluster, a minimum of four DSA modules must be installed and initialized as covered in the next chapter.

Chapter 3. Installing DSA Modules

After the SDN VE Controller is installed as described in the previous chapter, the Distributed Services Appliance (DSA) modules must be installed.

DSA modules are versatile software modules capable of being differentiated after installation to provide one of two vital functions in the SDN VE system:

- Distributed Connectivity Service (DCS)
Each DCS contains network information pertaining to nearby VMs, gateways and virtual switches in the SDN VE system. Tenant information is synchronized among partner modules to provide distributed virtual networking capabilities.
A minimum of two (2) DCS modules (installed on different hosts) are required for high-availability (HA) resilience. Three (3) are recommended.
- DOVE Gateway (DGW)
Each DGW can serve as a connection to an external, non-virtual network.
 - External Gateways are associated with a specific port in the physical network.
 - VLAN Gateways are associated with legacy VLAN broadcast domains.A minimum of two (2) DGW modules (installed on different hosts) are required for HA resilience.

Each VM used as an SDN VE entity must have a minimum allocation of 2 GB of memory.

The remainder of this chapter describes installing and initializing the DSA modules required for HA resilience.

Deploying the DSA Software

Though deploying DSA software can be accomplished using either the VMware vSphere Client, vSphere Web Client, or OVF Tool, the procedure shown in this *User Guide* depicts only the vSphere Client. If using one of the other tools, extrapolate from the information provided.

Follow these steps to deploy and start the required DSA modules:

1. Download the DSA OVA file from IBM.
2. Place the OVA file on a system that has access to the VMware vSphere Client (such as an administrative laptop).

Note: At least four DSA modules are required for high-availability resilience: A minimum of two (installed on different hosts) for DCS modules, and a minimum of two (installed on different hosts) for DGW modules. More can be installed if desired. Perform the remaining steps once for each module.

Install Using OVF Tool

Use the following command to download and install the OVA file:

```
$/opt/vmware/ovftool/ovftool --name=<NAME> [--powerOn]
--datastore=<"NAME"> --network="<Network>" dsa.ova vi://<vCenter
IP>/<DC>/Host?ip=<Host IP>
```

Replace the variables in the command with appropriate values as follows

Table 3. Command Parameters

Option	Description
Name	Name of the VM. For example: DSA1
powerOn	Specify this argument if want the SDN-VE Controller to power on when the command is executed.
Datastore	
Network	Select the VM network.

You will be prompted to enter the target host login credentials.

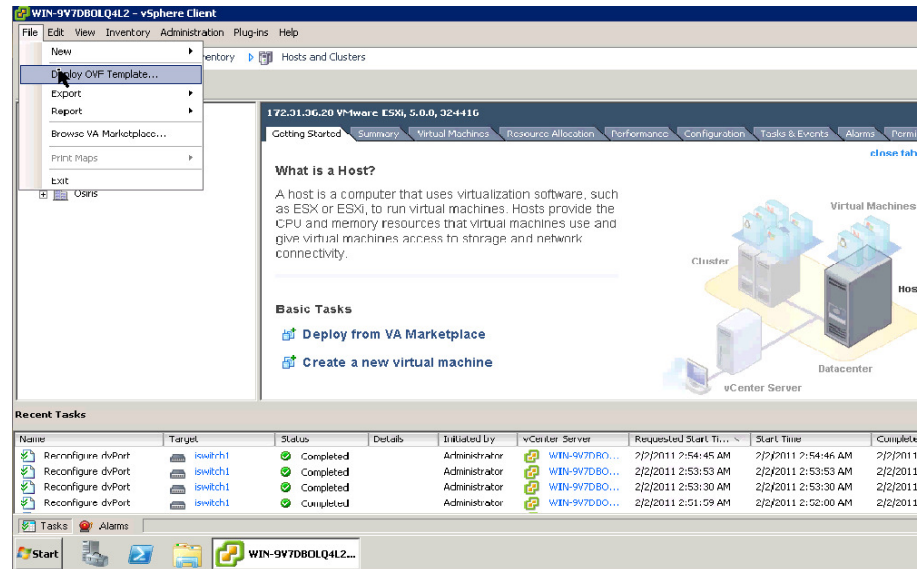
If you want to use vSphere Client to power on the DSA module, log in to the vSphere Client and follow [Step 11](#) onwards on [page 55](#).

Install Using VMware vSphere Client

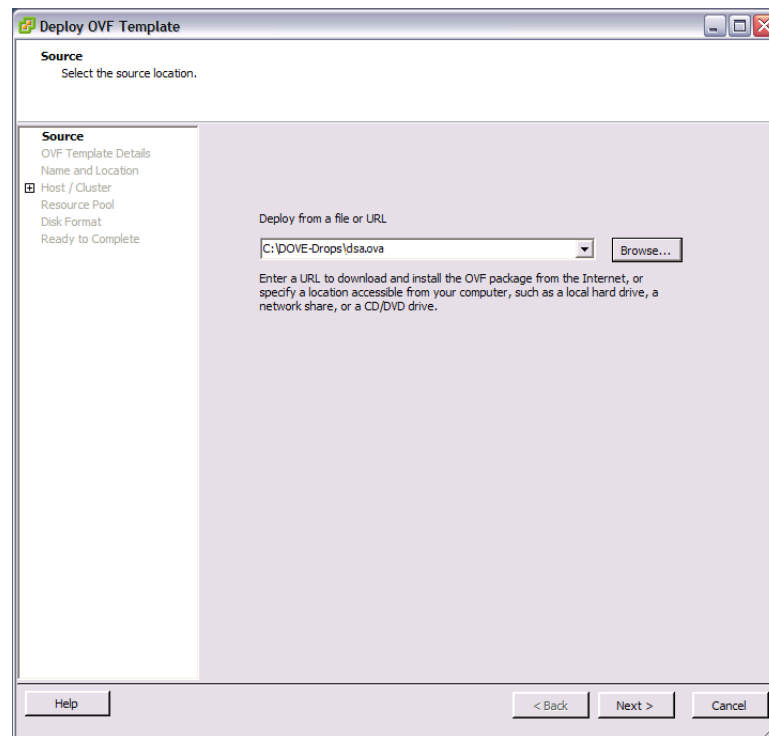
Follow these steps to install the SDN-VE DSA module:

1. Select an ESX host on which to deploy the DSA.
Each DSA is required to have Layer 3 connectivity to the designated vCenter and participating DMC modules.
2. Launch the VMware vSphere Client and connect either to the vCenter that manages the host where the DSA will be deployed or directly to the ESX host.

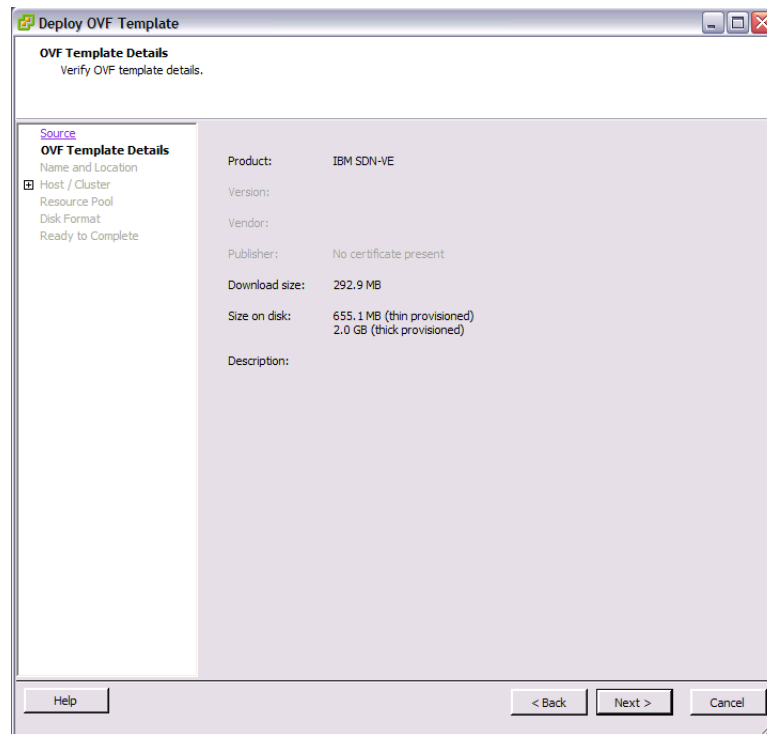
3. From the vSphere Client, select the target ESX host and choose **File > Deploy OVF Template** as shown below:




4. Select the location where the OVA file is stored and click **Next**.




5. Verify the OVA details and click **Next**.



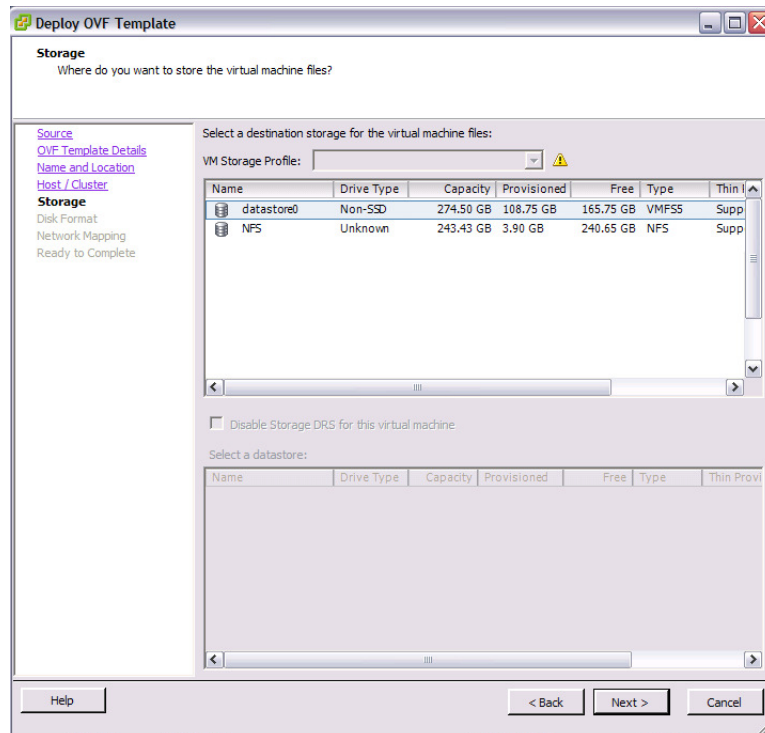
6. Provide a name for the DSA module and click Next.



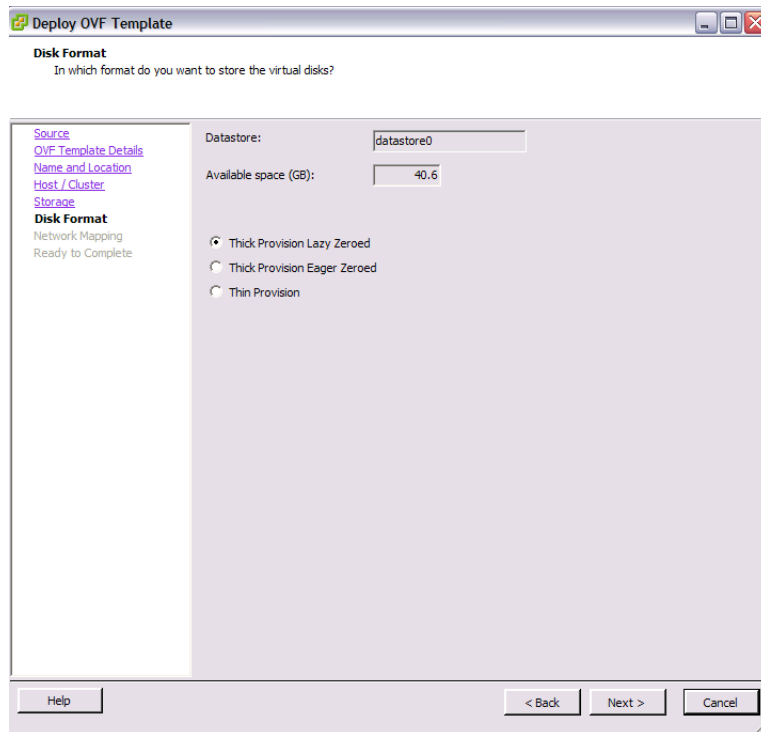
7. Specify the host or cluster on which to deploy the DSA and click **Next**.



8. Specify a location on the VM where DSA files should be stored, and click **Next**.

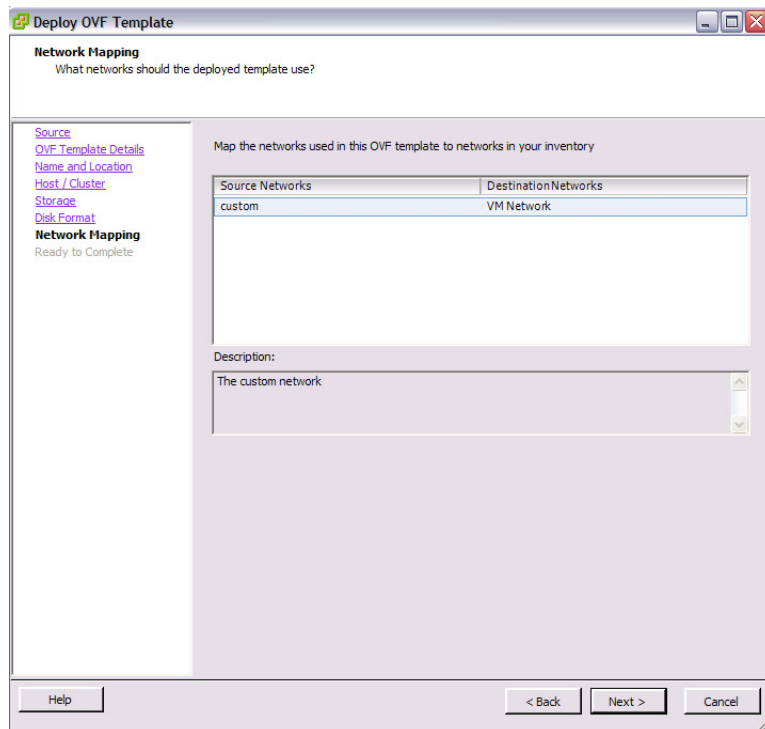


9. Select a disk format and click Next. The recommended format is Thick Provisioned Lazy Zeroed.



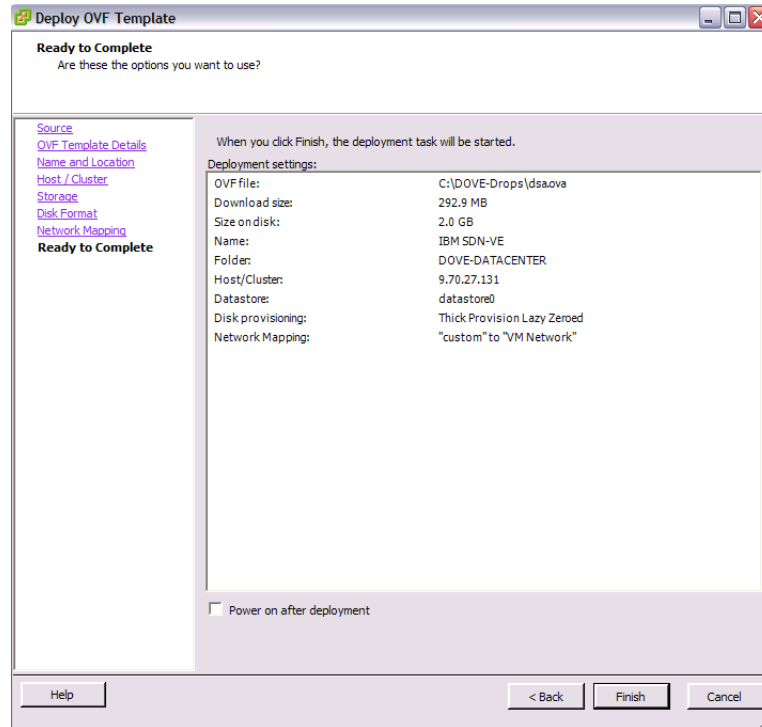
The screenshot shows the 'Deploy OVF Template' window with the 'Disk Format' tab selected. The window title is 'Deploy OVF Template'. The sub-header is 'Disk Format' with the question 'In which format do you want to store the virtual disks?'. On the left, there is a navigation pane with links: 'Source', 'OVF Template Details', 'Name and Location', 'Host / Cluster', 'Storage', 'Disk Format' (selected), 'Network Mapping', and 'Ready to Complete'. The main area shows 'Datastore:' as 'datastore0' and 'Available space (GB):' as '40.6'. There are three radio button options: 'Thick Provision Lazy Zeroed' (selected), 'Thick Provision Eager Zeroed', and 'Thin Provision'. At the bottom, there are buttons for 'Help', '< Back', 'Next >', and 'Cancel'.

10. Map the network for DSA controller use and click **Next**.




The screenshot shows the 'Deploy OVF Template' window with the 'Network Mapping' tab selected. The window title is 'Deploy OVF Template'. The sub-header is 'Network Mapping' with the question 'What networks should the deployed template use?'. On the left, there is a navigation pane with links: 'Source', 'OVF Template Details', 'Name and Location', 'Host / Cluster', 'Storage', 'Disk Format', 'Network Mapping' (selected), and 'Ready to Complete'. The main area has the instruction 'Map the networks used in this OVF template to networks in your inventory'. It features a table with two columns: 'Source Networks' and 'Destination Networks'. The table contains one row: 'custom' under 'Source Networks' and 'VM Network' under 'Destination Networks'. Below the table is a 'Description:' field with the text 'The custom network'. At the bottom, there are buttons for 'Help', '< Back', 'Next >', and 'Cancel'.

11. Verify the specified options, select the “Power on after deployment” option, and click **Next**.



This will initiate the DSA module VM deployment:



The DSA VM will power on when deployment is complete, and the DSA console will appear.

Initial DSA Setup

A minimum of four DSA modules are required. After installing the DSA modules on VM hosts, each DSA must be manually configured by entering commands into the built-in Command-Line Interface (CLI). Initially, the CLI can be accessed only through each DSA VM console on the vSphere Client. Later, if desired, the CLI can be accessed via remote SSH connection.

Note: Configuration of the DSA must be performed solely from the DSA console, and not through the vCenter interface (even in cases where the vCenter interface seems to allow it). However, host operations (such as adding ESX hosts and uplinks or assigning VM network interfaces to vDS ports or profiles) must be done through the vCenter interface.


Perform the following initial DSA setup for all DSA modules.

Start the DSA Module

When following the provided installation instructions, the DSA module automatically starts when the VM is powered on.

However, to manually access the console under other conditions, use the following procedure:

1. Log-in to the VMware vCenter via your vSphere Client.
2. Right-click on the target DSA VM and select the option to “Open Console.” Alternately, you can click on the Console icon.



The VM console for the selected DSA will appear.

Log In to the DSA

Note: You can use the GUI to specify the settings required to complete the installation. The use of CLI is documented in this *User Guide*.

Access the CLI using a remote SSH connection, or via virt-manager.

Once you are connected to the DSA, you are prompted to enter a login name and password. The default log-in user name and password are as follows:

Default user name: **admin**

Default password: **admin**

It is recommended that you change all default system password after initial configuration and as regularly as required under your network security policies.

Enter Global Configuration Mode

The DSA uses a CLI command set with multiple command modes. For an overview of CLI modes and features, see [“Command Basics” on page 187](#). The remainder of this chapter will display all commands necessary for initial configuration, but only those command relevant to the specific configuration tasks will be called into example.

After logging in, perform the following commands to enter the CLI Global Configuration mode:

```
SDN-VE-DSA> enable
SDN-VE-DSA# configure terminal
SDN-VE-DSA(config)#
```

Configure the DSA IPv4 Address (Optional)

Note: This section is not required if you are using DHCP.

Each DSA must have IPv4 connectivity to the SDN VE Controller modules that will participate in the SDN VE system.

By default, the DSA is enabled for dynamic IPv4 addressing using DHCP. If there is a DHCP server available in your network, the DSA will automatically acquire IPv4 address.

However, if DHCP is not available in your network or if you wish to override DHCP and configure static IPv4 addresses for the DSA, enter the following commands, depending on whether you prefer IPv4 address/netmask or CIDR notation.

If using static IP configuration, set the DSA IPv4 address using the following command

Using IPv4 Address and Nexthop using dotted-decimal (*a.b.c.d*):

```
SDN-VE-DSA(config)# ipmgmt set ip addr <DSA IPv4 address> mask <netmask>
SDN-VE-DSA(config)# ipmgmt set nexthop ip <IPv4 address>
```

(OR)

Using CIDR Notation dotted-decimal (*a.b.c.d/e*):

```
SDN-VE-DSA(config)# ipmgmt set ip cidr <DSA CIDR address>
```

You can verify DSA IPv4 address using the following command:

```
SDN-VE-DSA(config)# show ipmgmt

Mgmt IPv4: 9.121.62.42
Mask: 255.255.254.0
Nexthop: 9.121.62.1
```

Using DHCP

DHCP is used by default. However, if you have configured static IPv4 addresses and prefer to return to DHCP operation, enter the following command:

```
SDN-VE-DSA(config)# ipmgmt set dhcp
```

Note: Switching to DHCP will clear the static IPv4 addresses for the DSA.

Attach to the SDN VE Controller Cluster IPv4 Address

All DSA modules get the remainder of their functional configuration through the active SDN VE Controller operating at the SDN VE Controller cluster's HA external IPv4 address.

Use the DSA CLI to attach the DSA to the SDN VE Controller cluster. For each DSA, specify the SDN VE Controller cluster address (see [“Configure SDN VE DOVE HA” on page 45](#)) using the following Global Configuration command:

```
SDN-VE-DSA(config)# dmc set ip addr <DMC HA external IPv4 address>
```

Note: Be sure to use the HA external IPv4 address for the SDN VE Controller cluster, and not the individual primary or secondary SDN VE Controller IPv4 address. This helps preserve DSA communication resilience to the SDN VE Controller cluster in case the primary SDN VE Controller fails.

To verify DSA to SDN VE Controller connectivity, access the SDN VE Controller CLI and use the `show service-appliance` command:

```
SDN-VE@SDN-VE-Controller> configure terminal
SDN-VE@SDN-VE-Controller(config)# sdnve-dove terminal
```

```
SDN-VE-Controller(config-sdnve-dove)# show sdnve-dove service-appliances
```

DCS Service Appliances:

ID	IP	SERVICE CAPABILITY	ROLE ASSIGNED	AGE_TIME	CONFIG VERSION	BUILT VERSION
1	9.70.27.54	CS	N	13 s	0/ 1	1.0.0.130530
2	9.70.27.155	CS	N	15 s	0/ 1	1.0.0.130530
3	9.70.27.145	CS	N	11 s	0/ 1	1.0.0.130530
4	9.70.27.160	CS	N	12 s	0/ 1	1.0.0.130530

GW Service Appliances:

ID	IP	SERVICE CAPABILITY	ROLE ASSIGNED	AGE_TIME	CONFIG VERSION	BUILT VERSION
1	9.70.27.54	GW	N	7 s	0/ 1	1.0.0.130530
2	9.70.27.155	GW	N	15 s	0/ 1	1.0.0.130530
3	9.70.27.145	GW	N	0 s	0/ 1	1.0.0.130530
4	9.70.27.160	GW	N	15 s	0/ 1	1.0.0.130530

Note: Each of the installed DSA modules is shown in both the DCS list and the DGW list, but their roles as DCS or DGW is not yet assigned (ROLE ASSIGNED = N).

Specify DSA Roles

You can set each DSA to operate in either a DCS role or a DGW role. These roles are mutually exclusive: At any given time, the DSA can operate in one or the other, but not both. Roles are defined using the SDN VE Controller CLI (not via the DSA itself).

On the SDN VE Controller, assign DCS roles to at least two unassigned modules (on different hosts) using the following Global Configuration mode command:

```
SDN-VE-Controller(config-sdnve-dove)# service role dcs ids <list of target DSA modules>
```

Example:

```
SDN-VE-Controller(config-sdnve-dove)# service role dcs ids 1,2
```

where the list is a comma separated list of numeric DSA IDs as seen in the `show service-appliance` command (see [page 58](#)).

Also assign DGW roles for two unassigned modules (on different hosts) using the similar command:

```
SDN-VE-Controller(config-sdnve-dove)# service role dgw ids <list of target DGW modules>
```

Example:

```
SDN-VE-Controller(config-sdnve-dove)# service role dgw ids 3,4
```

Verify the settings using the `show service-appliance` command:

```
SDN-VE-Controller(config-sdnve-dove)# show sdnve-dove service-appliances
```

DCS Service Appliances:							
ID	IP	SERVICE CAPABILITY	ROLE ASSIGNED	AGE_TIME	CONFIG VERSION	BUILT VERSION	
1	9.70.27.54	CS	Y	13 s	0/ 1	1.0.0.130530	
2	9.70.27.155	CS	Y	15 s	0/ 1	1.0.0.130530	
3	9.70.27.145	CS	N	11 s	0/ 1	1.0.0.130530	
4	9.70.27.160	CS	N	12 s	0/ 1	1.0.0.130530	

GW Service Appliances:							
ID	IP	SERVICE CAPABILITY	ROLE ASSIGNED	AGE_TIME	CONFIG VERSION	BUILT VERSION	
1	9.70.27.54	GW	N	7 s	0/ 1	1.0.0.130530	
2	9.70.27.155	GW	N	15 s	0/ 1	1.0.0.130530	
3	9.70.27.145	GW	Y	0 s	0/ 1	1.0.0.130530	
4	9.70.27.160	GW	Y	15 s	0/ 1	1.0.0.130530	

Once roles are successfully set, the “ROLE ASSIGNED” field will be Y (Yes) in the appropriate role table.

Configure Tunnel Endpoints

For DSAs configured with the `dgw` role, you must configure tunnel endpoint IP address so they can receive packets from the DOVE virtual switches. See [“Installing the SDN VE vSwitch” on page 47](#).

Following are the steps for configuring the tunnel endpoint:

At the SDN VE Controller CLI prompt:

1. Configure service gateway IP address:

```
SDN-VE-Controller(config-sdnve-dove)# service dgw id <dgw ID>
add-interface ip <tunnel IP address> mask <netmask> nexthop <nexthop IP address>
dovetunnel
```

2. Verify the IP address at the controller:

```
SDN-VE-Controller(config-sdnve-dove)# show sdnve-dove dgw-interfaces id
<dgw ID>
```

DGW Index 2:
DGW IPv4 Stats:

ID	GWIDX	IP	MASK	NEXTHOP	TYPE	VLAN
--	-----	--	----	-----	----	----
1	2	2.2.2.31	255.255.255.0	2.2.2.1	dovetunnel	0

At the DSA CLI prompt:

3. Access the DSA configuration mode on the DSA module:

```
SDN-VE-DSA> enable
SDN-VE-DSA# configure terminal
SDN-VE-DSA(config)#
```

4. Verify the IP address at the gateway:

```
SDN-VE-DSA(config)# show ipv4-interfaces
0: 127.0.0.1
1: 9.121.62.31
2: 1.1.1.31
```

Next Steps

Once DSA roles have been assigned for all required modules, a DS 5000V virtual switch must be installed and initialized as covered in the next chapter.

Chapter 4. Installing the SDN VE 5000V Distributed vSwitch

The IBM SDN VE 5000V Distributed vSwitch (5000V), version 1.2, is a virtual distributed switch (vDS) solution for VMware. It provides network switching within the SDN VE network fabric.

This chapter describes installing the 5000V as part of the IBM SDN VE solution. These steps vary from those stated in the IBM System Networking DS 5000V User Guide, which covers installing the 5000V as a stand-alone vDS (without SDN VE).

Deploying the 5000V Controller Software

The 5000V controller software can be deployed using either the VMware vSphere Client, vSphere Web Client, or OVF Tool. The procedure shown in this *User Guide* depicts the OVF tool and vSphere Client. If using vSphere Web Client, extrapolate from the information provided.

Follow these steps to deploy and start the required 5000V controller software:

1. Download the 5000V version 1.2 controller OVA file from IBM.
2. Place the OVA file on a system that has access to the VMware vSphere Client (such as an administrative laptop).

Install Using OVF Tool

Use the following command to download and install the OVA file:

```
$ /opt/vmware/ovftool/ovftool --name=<NAME> [--powerOn]
--datastore=<"NAME"> --network="<Network>" 5000v-controller.ova
vi: //<vCenter IP>/<DC>/Host?ip=<Host IP>
```

Replace the variables in the command with appropriate values as follows

Table 4. Command Parameters

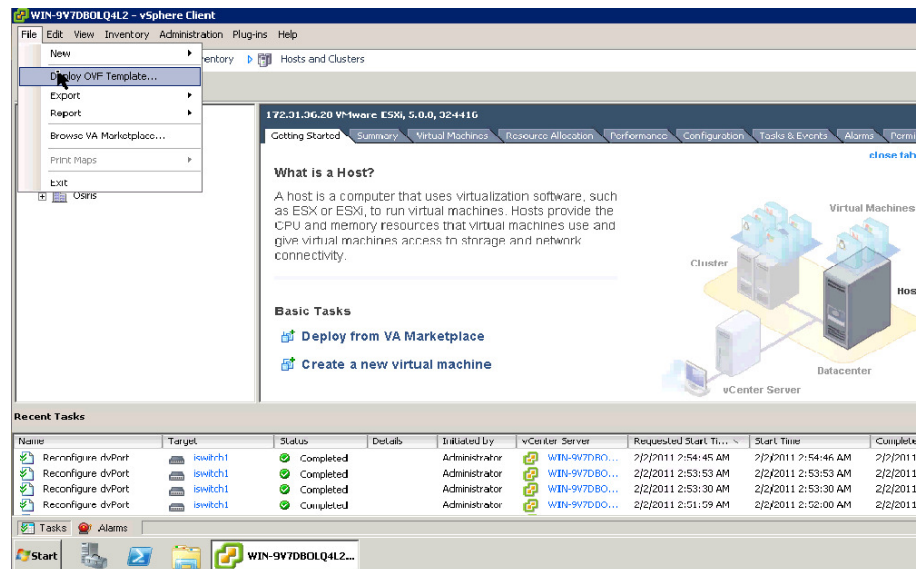
Option	Description
Name	Name of the VM. For example: vDS1
powerOn	Specify this argument if want the SDN-VE Controller to power on when the command is executed.
Datastore	
Network	Select the VM network.

You will be prompted to enter the target host login credentials.

If you want to use vSphere Client to power on the 5000V Controller, log in to the vSphere Client and follow [Step 11](#) onwards on [page 66](#).

Install Using VMware vSphere Client

1. Specify an ESX host on which to deploy the controller.
The controller host merely provides an environment in which the 5000V controller appliance will run. It is not required to participate as a vDS host and may be a different class of device than those where the vDS host modules will later be installed. The primary requirement is for the controller host to have Layer 3 connectivity to the designated vCenter and the SDN VE cluster.
2. Launch the VMware vSphere Client and connect either to the vCenter that manages the host where the 5000V controller will be deployed or directly to the ESX host.
3. From the vSphere Client, select the target ESX host and choose **File > Deploy OVF Template** as shown below:



4. Select the location where the OVA file is stored and click **Next**.

The dialog box is titled "Deploy OVF Template". The main heading is "Source" with the instruction "Select the source location." On the left, a tree view shows "Source" selected, with sub-items: "OVF Template Details", "Name and Location", "Host / Cluster", "Resource Pool", "Disk Format", and "Ready to Complete". The main area is titled "Deploy from a file or URL" and contains a text box with the path "C:\DOVE-Drops\5000V-Controller.ova" and a "Browse..." button. Below this, a note states: "Enter a URL to download and install the OVF package from the Internet, or specify a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive." At the bottom, there are buttons for "Help", "< Back", "Next >", and "Cancel".


5. Verify the OVA details and click **Next**.

The dialog box is titled "Deploy OVF Template". The main heading is "OVF Template Details" with the instruction "Verify OVF template details." On the left, the tree view shows "OVF Template Details" selected, with sub-items: "Source", "Name and Location", "Host / Cluster", "Resource Pool", "Disk Format", and "Ready to Complete". The main area displays the following details:


Product:	IBM SN DS 5000V
Version:	
Vendor:	
Publisher:	No certificate present
Download size:	74.1 MB
Size on disk:	217.3 MB (thin provisioned) 900.0 MB (thick provisioned)
Description:	

At the bottom, there are buttons for "Help", "< Back", "Next >", and "Cancel".

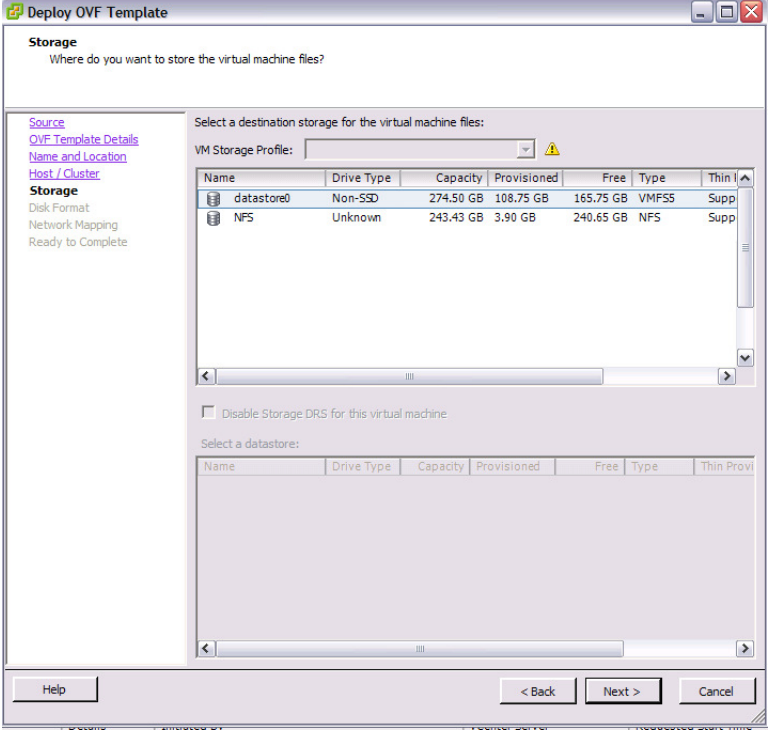
6. Provide a name the 5000V controller and click Next.



7. Specify the host or cluster on which to deploy the 5000V controller and click **Next**.



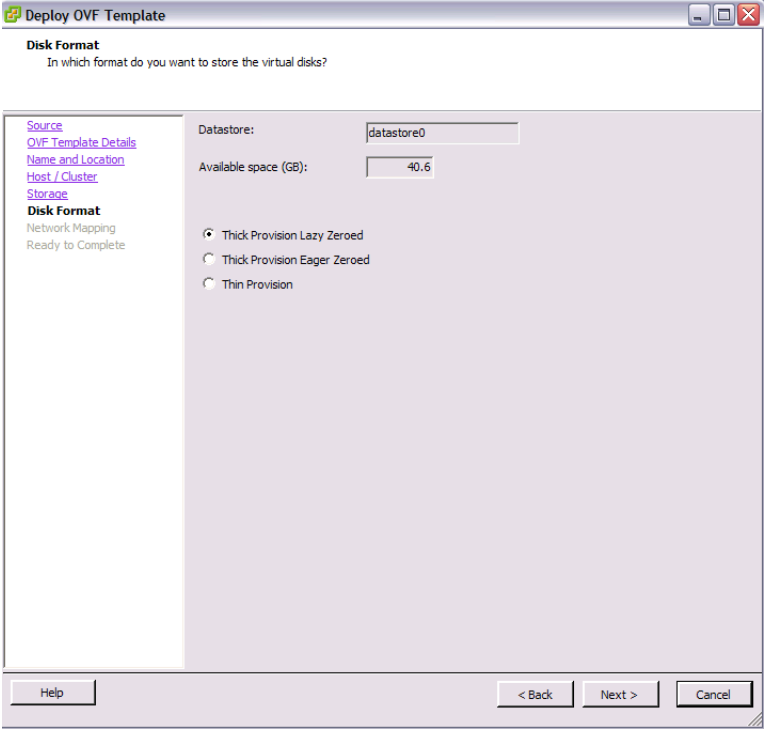
- Specify a location on the VM where 5000V controller files should be stored, and click **Next**.



The screenshot shows the 'Storage' step of the 'Deploy OVF Template' wizard. The title bar reads 'Deploy OVF Template'. The main heading is 'Storage' with the subtext 'Where do you want to store the virtual machine files?'. On the left, a navigation pane lists 'Source', 'OVF Template Details', 'Name and Location', 'Host / Cluster', 'Storage' (selected), 'Disk Format', 'Network Mapping', and 'Ready to Complete'. The main area is titled 'Select a destination storage for the virtual machine files:'. It includes a 'VM Storage Profile:' dropdown menu with a warning icon. Below is a table with columns: Name, Drive Type, Capacity, Provisioned, Free, Type, and Thin Provisioning. The table contains two rows: 'datastore0' (Non-SSD, 274.50 GB, 108.75 GB, 165.75 GB, VMFS5, Supp) and 'NFS' (Unknown, 243.43 GB, 3.90 GB, 240.65 GB, NFS, Supp). Below the table is a checkbox 'Disable Storage DRS for this virtual machine'. At the bottom are 'Help', '< Back', 'Next >', and 'Cancel' buttons.

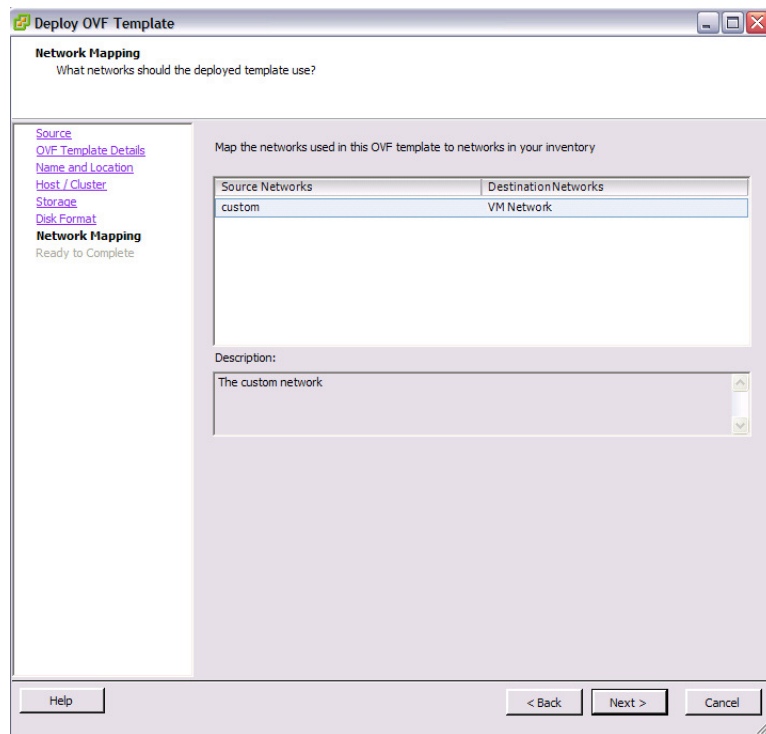
Name	Drive Type	Capacity	Provisioned	Free	Type	Thin Provisioning
datastore0	Non-SSD	274.50 GB	108.75 GB	165.75 GB	VMFS5	Supp
NFS	Unknown	243.43 GB	3.90 GB	240.65 GB	NFS	Supp

- Select a disk format and click Next. The recommended format is Thick Provisioned Lazy Zeroed.

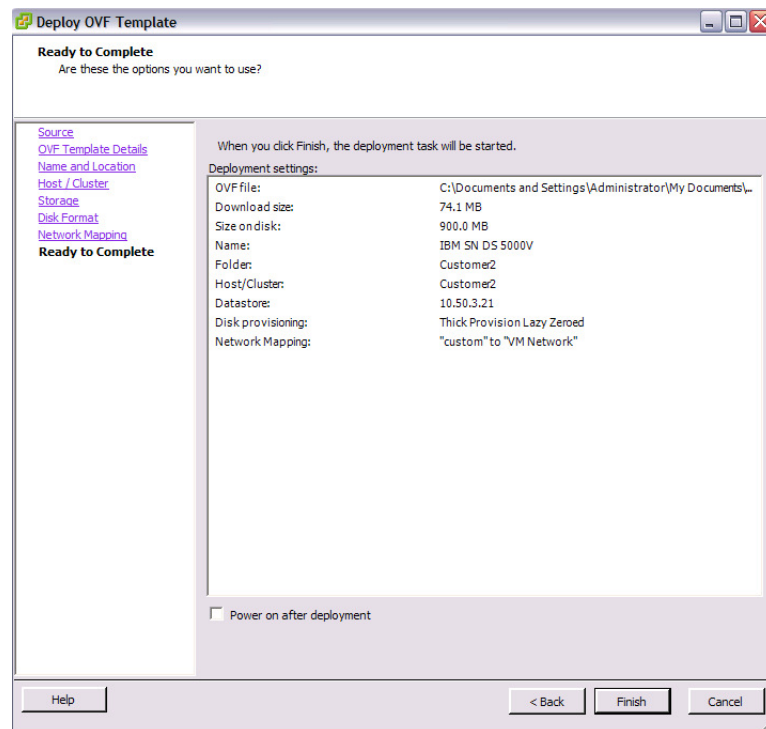


The screenshot shows the 'Disk Format' step of the 'Deploy OVF Template' wizard. The title bar reads 'Deploy OVF Template'. The main heading is 'Disk Format' with the subtext 'In which format do you want to store the virtual disks?'. On the left, a navigation pane lists 'Source', 'OVF Template Details', 'Name and Location', 'Host / Cluster', 'Storage', 'Disk Format' (selected), 'Network Mapping', and 'Ready to Complete'. The main area shows 'Datastore:' as 'datastore0' and 'Available space (GB):' as '40.6'. Below are three radio button options: 'Thick Provision Lazy Zeroed' (selected), 'Thick Provision Eager Zeroed', and 'Thin Provision'. At the bottom are 'Help', '< Back', 'Next >', and 'Cancel' buttons.


10. Map the network for 5000V controller use and click **Next**.



11. Verify the specified options, select the “Power on after deployment” option, and click **Next**.



This will initiate the 5000V controller VM deployment:



The 5000V controller VM will power on when deployment is complete, and the controller VM console will appear.

Initial 5000V Controller Setup

The 5000V must be manually configured by entering commands into the controller's built-in Command-Line Interface (CLI). Initially, the CLI can be accessed only through the 5000V controller VM console on the vSphere Client. Later, if desired, the CLI can be accessed via remote Telnet or SSH connections.


Note: Configuration of the 5000V vDS must be performed solely from the 5000V CLI, and not through the vCenter interface (even in cases where the vCenter interface seems to allow it). However, host operations (such as adding ESX hosts and uplinks or assigning VM network interfaces to vDS ports or profiles) must be done through the vCenter interface.

Start the 5000V Controller

When following the provided installation instructions (see [Step 11](#) on [page 66](#)), the controller console automatically appears when the 5000V controller VM is powered on.

However, to manually access the controller console under other conditions, use the following procedure:

1. Log-in to the VMware vCenter via your vSphere Client.
2. Right-click on the 5000V controller VM and select the option to "Open Console." Alternately, you can click on the Console icon.



The VM console for the 5000V controller will appear.

Set the language

When the SDN-VE Controller CLI opens, you will be prompted to set the language:

Please select a language and press enter (eg. 0 for English):

Examine the License Agreement

The first time the 5000V controller is started, you will be prompted to read the Software Licence Agreement. When you select a language, the SLA will be displayed.

When you are finished examining the SLA, select **1** if you wish to accept the terms.

If you accept the SLA, the 5000V controller login prompt will appear.

Log In to the 5000V Controller

CLI access is controlled through the use of a login name and password. Once you are connected to the 5000V controller, you are prompted to enter a login name and password. The default log-in user name and password are as follows:

Default user name: **admin**

Default password: **admin**

It is recommended that you change all default system password after initial configuration and as regularly as required under your network security policies.

Enter Global Configuration Mode

The 5000V controller uses a rich CLI command set with multiple command modes. For an overview of CLI modes and features, refer to the *Distributed Switch 5000V User Guide*. The remainder of this chapter will display all commands necessary for initial configuration, but only those command relevant to the specific configuration tasks will be called into example.

After logging in, perform the following commands to enter the CLI Global Configuration mode:

```
5000V> ena
5000V# configure terminal
5000V(config)#
```

The `ena` command initiates executive privilege mode, and the `configure terminal` command readies the controller for configuration.

Verify the 5000V Controller Version

The SDN VE solution requires version 1.2 of the DS 5000V. To verify the correct version of software has been deployed, use the following command:

```
5000V(config)# show running-config
```

Near the top of the output, a “Software Version” message is displayed. Verify that the version number is `1.2.0` or higher.

If an earlier version is deployed, refer to “Updating the Switch Software Image” in the “Boot Options” chapter in the *Distributed Switch 5000V User Guide*.

Configure the 5000V IPv4 Addresses (Optional)

The 5000V controller must have IPv4 connectivity to the VMware vCenter, as well as the hosts that will participate in the SDN VE system.

By default, the 5000V controller is enabled for dynamic IPv4 addressing using DHCP. If there is a DHCP server available in your network, the controller will automatically acquire its IPv4 address and gateway configuration. If using DHCP, you can skip static address configuration.

However, if DHCP is not available in your network or if you wish to override DHCP and configure static IPv4 addresses for the 5000V controller, enter the following command:

```
5000v(config)# interface ip-mgmt address <IPv4 address> [<mask>]
```

where *IPv4 address* is the address of the controller in dotted-decimal notation, optionally followed by the network *mask* used for creating an address range

If desired, you can also configure the *gateway* IPv4 address that the controller should use for outbound traffic:

```
5000v(config)# interface ip-mgmt gateway <gateway IPv4 address>  
5000v(config)# interface ip-mgmt gateway enable
```

Create the Global vDS Instance

The 5000V controller must be associated with a virtual distributed switch (vDS) for a particular virtual data center. The following CLI commands on the controller VM console are used to create the required association to the vCenter:

```
5000v(config)# iswitch vcenter <vCenter IPv4 address> <user name>
```

The *vCenter IPv4 address* represents the vCenter to which the 5000V will connect and *username* is the vCenter login name.

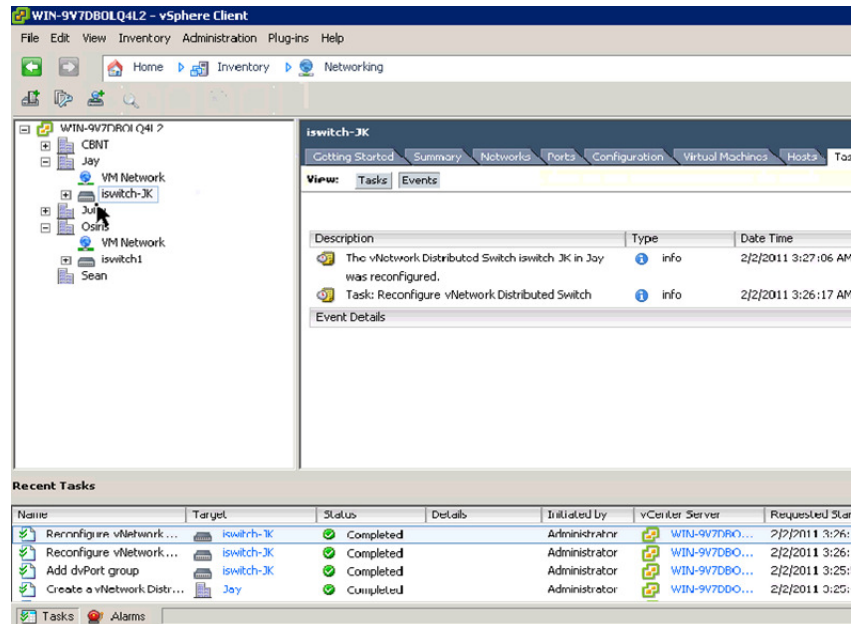
The system will then prompt you for the vCenter login password and its logical port number. By default, the vCenter operates on recommended TCP port number 443. However, if your vCenter communicates on a different port, enter the port number configured for the service.

Next, the 5000V controller must be associated with the vDS:

```
5000v(config)# iswitch vds <vDS name> <datacenter name>
```

Note: The assigned names cannot include internal spaces.

When this configuration is complete, the 5000V vDS will appear at the vCenter in the **Home > Inventory > Networking** view:



Note: Once the controller is associated a vDS in the vCenter, whenever the IPv4 address of the 5000V controller is changed (statically or via DHCP renewal), you must save the 5000V configuration and reload the controller in order to reestablish the required association.

Attach to the DMC Module Cluster IPv4 Address

The 5000V must coordinate its ongoing configuration with the active DMC Controller module. Use the 5000V CLI to select the DMC module cluster's HA external IPv4 address (from ["Establish Unified Controller High-Availability" on page 45](#)). Specify the address using the following Global Configuration command:

```
5000V(config)# iswitch dmc <SDN-VE HA external IPv4 address>
```

Note: Be sure to use the HA external IPv4 address for the DMC module cluster, and not the individual primary or secondary DMC module IPv4 address. This helps preserve 5000V communication with the DMC module cluster in case the primary DMC module fails.

Verify the DMC module configuration using the `show running-config` command on the 5000V controller CLI and examining the `iswitch` output elements.

Also verify that a DOVE Tunnel End Point (TEP) profile has been automatically created at the vCenter.

Next Steps

Once installation and initial setup of the DMC module, DSA, and 5000V elements are complete, the system is ready for virtual network configuration as discussed in the next chapter of this *User Guide*.

Chapter 5. Virtual Network Configuration

Overview

This section provides an overview of the SDN setup including the network layers and the key configurations.

The SDN VE setup has three layers:

Management Network: All management between the hosts and the SDN VE components rides over this layer. The layer resides on the virtual switch in each host, and has an uplink to the physical environment.

Tunnel Endpoint (TEP) Network: The TEP network resides between the overlay and underlay networks. Each host must be configured with an IP address to enable communication with the TEP network.

Overlay Network: A virtual network that can be defined as a tenant with multiple subnets. Each host communicates with the overlay network via the virtual switch.

To enable host-to-host communication, gateways are required. These are defined in the Unified Controller.

To enable VM-to-VM communication across tenants, Distributed External Gateways are required.

Before proceeding with the configuration, ensure you have:

- Installed the Unified Controller
- Installed the Distributed Services Appliance (DSA)
- Defined the Distributed Connectivity Service (DCS) and Distributed Gateways (DGW)
- Installed the SDN VE 5000V Distributed vSwitch and the 5000V Host module .

Following is a summary of the configuration procedure. The details are provided later in this chapter.

- Define the overlay network on the Unified Controller:
 - Create tenants.
 - Create connectivity groups within each domain.
 - Create IP subnets and assign them to the networks.
 - Define connectivity group policies.
 - Export connectivity groups to the virtual switch.
- Configure Distributed VLAN Gateway and Distributed External Gateway.
- Install the 5000V Host module
- Define the underlay network.
- Attach hosts to the vDS.
- Configure TEP address.
- Assign each VM to the appropriate port group.

Overlay Configuration

Once basic installation and initial configuration is complete, the overlay network can be configured. The overlay network consists of tenants, connectivity groups, the address spaces that will be mapped to the connectivity groups, and the policies between connectivity groups. Overlay configuration is performed via the Unified Controller module.

Create Tenants

Tenants are created from the SDN VE global configuration mode using the following command:

```
tenant add id <tenant_id> name <tenant_name> type dove
[descr <description>]
```

For Example:

```
SDN-VE @SDN-VE-Controller(config)# tenant add id 3 name Corp type dove
descr Corporate
Tenant created with UUID = 3
```

To add a replication factor, use the following command:

```
tenant update id <id> [name <tenant_name>] [repfactor <replication factor>]
[descr <description>]
```

For Example:

```
SDN-VE @SDN-VE-Controller(config)# tenant update id 3 repfactor 2
```

Create tenants are listed using the `show tenant` command:

```
SDN-VE @SDN-VE-Controller > show tenant
```

Id	Name	Domain_Type	Replication factor	Description
---	-----	-----	-----	-----
1	DOVE ADMIN	DOVE	2	Admin Tenant for DOVE, Created at startup
2	OF ADMIN	OF	0	Admin Tenant for OF, Created at startup
3	Corp	DOVE	2	Corporate

Create Connectivity Groups

Connectivity Groups are created from the SDN VE Tenant configuration mode. This mode can be accessed using the following command:

```
tenant id <tenant_id>
```

Connectivity group is created using the following command:

```
group add name <CG name> admin-state <CG status> [vnid <ID>]
[traffic <traffic type>] [precedence <level>] [limitDelay <value >]
[limitThroughput <value >] [limitReliability <value >]
[average_rate <value in KBps >] [peak_rate <value in KBps >]
[burst_rate <value in KiloBytes >] [id <group ID>]
[group-type {dedicated | shared | external}] [isNeutron {true | false}]
```


For Example:

```
SDN-VE @SDN-VE-Controller(config)# tenant id 3
SDN-VE @SDN-VE-Controller(config-tenant-3)# group add name Corp_HR
admin-state up vnid 11 id 3 group-type dedicated
```

Create groups are listed using the `show group` command:

```
SDN-VE @SDN-VE-Controller > show group

Id                : 3
Name              : Corp_HR
Vnid              : 11
Tenant_id         : 3
Admin State       : true
Status            : ACTIVE
Group_type        : dedicated
isNeutron         : false
Waypoint          : false
QoS :
  Traffic type    : BEST_EFFORT
  Precedence type : FLASH_OVERRIDE
  Limit delay     : 0
  Limit throughput : 0
  Limit reliability : 0
Rate Limits :
  Average_rate    : 0
  Peak_rate       : 0
  Burst_rate      : 0
```

Create Subnets

Create subnets from the Tenant Configuration mode using the following command:

```
subnet add name <Subnet name> cidr <CIDR IPv4 address> [id <Subnet ID>]
[isNeutron <true or false>] [subnet_type {dedicated | shared | external}]
[gateway <gateway IP address>]
[allocation_pools start <Starting IP address> end <Ending IP address>]
[vlan <VLAN ID>]
```

For example:

```
SDN-VE @SDN-VE-Controller (config-tenant-3)# subnet add name Corp_HR_Sub
cidr 10.1.1.0/24 id 10 type dedicated gateway 10.1.1.1
```

Created subnets are listed using the `show subnet` command:

```
SDN-VE @SDN-VE-Controller (config-tenant-3)# show subnet

Id                : 10
Tenant Id         : 3
Name              : Corp_HR_Sub
isNeutron         : false
CIDR              : 10.1.1.0/24
Subnet type       : dedicated
IP Version        : 4
Gateway Ip        : 10.1.1.1
Allocation pools  :
  Pool 1 : 10.1.1.1 - 10.1.1.254
```

Note: The *gateway* IPv4 address should be the default route for all endpoints that attach to the network to which this subnet is bound.

Bind Subnets to the Connectivity Group

Because a subnet can be bound to multiple connectivity groups, it is necessary to configure bindings through the Group Configuration mode. This mode is accessed via the Tenant Configuration mode, using the following command:

```
group id <group id>
```

where the group ID is as shown using the `show group` command.

For example:

```
SDN-VE @SDN-VE-Controller(config-tenant-3-group)# subnet attach id 10
```

Connectivity groups can contain subnets of only one type. The type that is added first determines what other subnets can be added to the network. So if a `dedicated` type subnet was added first, all subsequent subnets that are added to that connectivity group need to be of type `dedicated`.

Define Connectivity Group Policy (Optional)

Policies are defined in the Tenant Configuration mode using the following command:

```
cgpolicy add id <id1> id <id2> traffic-type <traffic type> directional  
<traffic direction>
```

Policies defined between connectivity groups enables communication between the groups.

For example (assume we have added another connectivity group with ID 4)

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# cgpolicy add id 3 id 4  
traffic-type Unicast directional BI_DIRECTIONAL
```

Export Networks to the SDN VE 5000V vSwitch

Before traffic can flow, you must export the created virtual networks to the 5000V vSwitch. This makes the information available to the vDS virtual switches for connected VMs.

To export a network, use the following Group Configuration command:

```
export ip <virtual switch IP address>
```

where `<virtual switch IP address>` is the IPv4 address of the 5000V vSwitch .

For example:

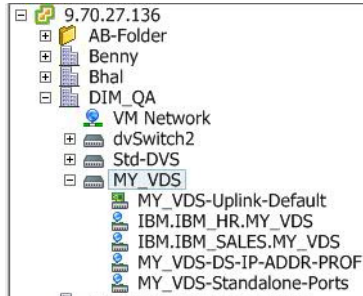
```
SDN-VE @SDN-VE-Controller(config-tenant-3-group)# export ip 9.121.62.27
```

On the 5000V vSwitch console, messages will appear when the network profiles are created. For example:

```
Jun  4 2013 18:43:29 5000V:SYSTEM-INFO: Saved configuration to flash
successfully!

Jun  4 2013 18:43:29 5000V:SYSTEM-ALERT: Profile [IBM.IBM.HR.MY_VDS] got
created from DMC, config saved
```

To verify that the profiles have been created, use either vCenter or the 5000V controller show running-config command. For example:



```
5000V(config)# show running-config

Building configuration...
#
#switch-type "IBM System Networking Distributed Switch 5000v"
#Software Version 1.1.0.130603
#!!!!DO NOT EDIT ANYTHING ABOVE THIS LINE!!!!
#
!
!
iswitch vcenter 8.70.27.136 root 0x559b5fe219e61dec 443
iswitch vds MY_VDS DIM_QA dvs-9609 datacenter-1686
iswitch dmc 9.70.27.245
iswitch doveprof IBM.IBM_HR.MY_VDS 10 141 dvportgroup-9639
!
iswitch doveprof IBM.IBM_SALES.MY_VDS 10 151 dvportgroup-9640
!
!
!
!
iswitch doveprof IBM.IBM_HR.MY_VDS dvportgroup-9639
        vnid 1
iswitch doveprof IBM.IBM_SALES.MY_VDS dvportgroup-9640
        vnid 2
!
end
```

Externalizing the Overlay Networks

To connect an overlay network to a traditional network, a gateway is used. There are two types of gateways: those that connect a virtual network to a legacy VLAN environment, and those that connect a virtual network to external hosts, including the Internet. A particular DSA assigned a role as a Distributed Gateway (DGW) can only function as one gateway type.

You can view the available gateways using the following command:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove
service-appliances
```

DCS Service Appliances:						
ID	IP	SERVICE CAPABILITY	ROLE ASSIGNED	AGE_TIME	CONFIG VERSION	BUILT VERSION
1	9.70.27.54	CS	N	13 s	0/ 1	1.0.0.130530
2	9.70.27.155	CS	N	15 s	0/ 1	1.0.0.130530
3	9.70.27.145	CS	N	11 s	0/ 1	1.0.0.130530
4	9.70.27.160	CS	N	12 s	0/ 1	1.0.0.130530

GW Service Appliances:						
ID	IP	SERVICE CAPABILITY	ROLE ASSIGNED	AGE_TIME	CONFIG VERSION	BUILT VERSION
1	9.70.27.54	GW	N	7 s	0/ 1	1.0.0.130530
2	9.70.27.155	GW	N	15 s	0/ 1	1.0.0.130530
3	9.70.27.145	GW	N	0 s	0/ 1	1.0.0.130530
4	9.70.27.160	GW	N	15 s	0/ 1	1.0.0.130530

Configure a VLAN Gateway

Ensure you have configured the tunnel endpoint IPv4 address. See [“Configure Tunnel Endpoints” on page 59](#).

If not yet configured, configure the tunnel endpoint's (TEP) IPv4 address with the following command:

```
service dgw id <Service appliance ID> add-interface ip <IPv4 address>
mask <netmask> nexthop <gateway IPv4> {dovetunnel|external}
vlan <VLAN ID>
```

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# service dgw id 1 add-interface
ip 2.2.2.23 mask 255.255.255.0 nexthop 2.2.2.254 dovetunnel vlan 0
```

After adding the TEP, use the Group Configuration mode to set the VLAN:

```
vlan-gateway add dgw_id <DGW ID> vlan <VLAN ID>
```

```
SDN-VE @SDN-VE-Controller(config-tenant-1-group)# vlan-gateway add dgw_id 1 vlan 201
```

For example, this instructs the gateway shown in the service appliance list as index #3 (with a management IPv4 address of 9.70.27.145) to map traffic on connectivity group 1 to VLAN ID 201.

This completes the VLAN Gateway setup on the SDN VE Controller module.

Configure an External Gateway

External gateway configuration is required for the VM network (data network) to communicate with the external network (For example: the Internet). You need two IPv4 addresses for external gateway configuration:

- Tunnel end point address for the data network. This address provides Layer 3 connectivity to the destination.
- External IPv4 address to connect with the external network.

Note: Only one External IPv4 address can be configured.

Following is the command sequence for configuring an external gateway (EGW):

1. Configure the external IPv4 address with the command:

```
service dgw id <Service appliance ID> add-interface ip <IPv4 address>
mask <netmask> nexthop <gateway IPv4> {dovetunnel|external}
[vlan <VLAN ID>]
```

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# service dgw id 4 add-interface
ip 7.7.7.24 mask 255.255.255.0 nexthop 7.7.7.100 external vlan 200
```

2. Enter gateway configuration mode:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# service gateway id 4
SDN-VE @SDN-VE-Controller (config-sdnve-dove-gateway)#
```

3. Specify forwarding rules:

```
fwd-rule add {vnid <VNID>|group_id <Group ID>}
overlayip <overlay IPv4> floating-ip <floating IPv4 address>
[proxy-min-ip <proxy start IP>] [proxy-max-ip <proxy end IP>]
[protocol <protocol>] [port <logical port>] [overlay-port <overlay port>]
```

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove-gateway)# fwd-rule add
group_id 1 overlay-ip 10.1.1.10 floating-ip 20.20.20.1 proxy-min-ip
2.2.2.100 proxy-max-ip 2.2.2.150 protocol 6 port 5001 overlayport 5001
SDN-VE @SDN-VE-Controller (config-sdnve-dove-gateway)# exit
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# exit
SDN-VE @SDN-VE-Controller (config)#
```

This command sets up addresses 2.2.2.100 and 2.2.2.150 on gateway index #4 (with a management IPv4 address of 9.70.27.160) as NAT addresses for devices attached to Group 1.

4. Configure a policy to be applied between two connectivity groups:

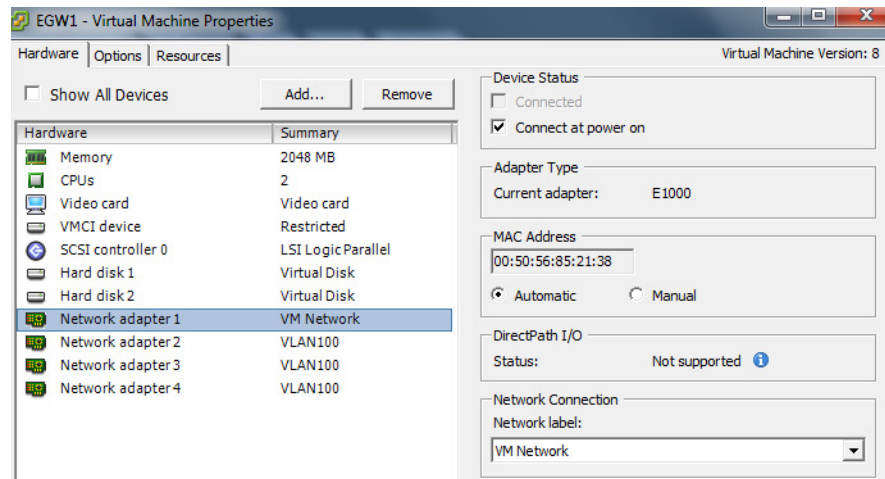
```
cgpolicy add id <id1> id <id2> traffic-type <traffic type>
directional <traffic direction> [snat start-ip <Start IP address> end-ip
<End IP address> start-port <Start port number> end-port <End Port
number>]
```

```
SDN-VE @SDN-VE-Controller (config)# tenant id 3
SDN-VE @SDN-VE-Controller (config-tenant-1)# cgpolicy add id 3 id 4
traffic-type UNICAST directional BI_DIRECTIONAL
```

This completes the configuration of external gateways on the SDN VE Controller.

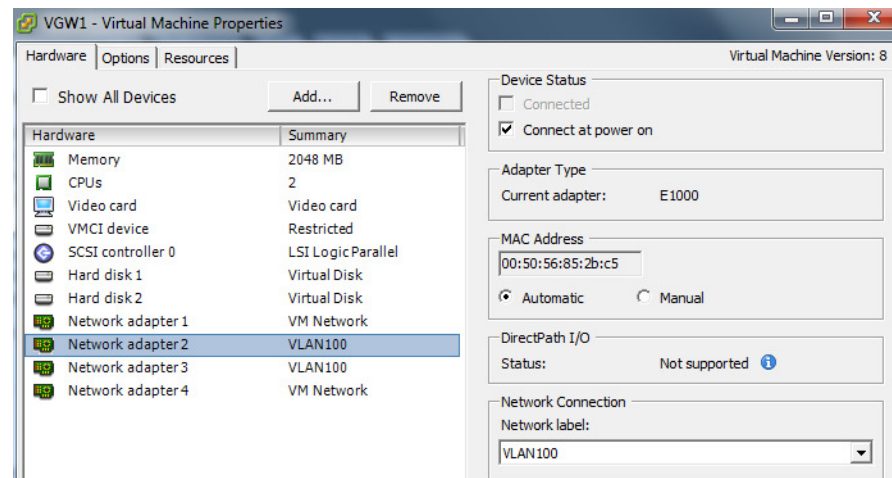
Configuration of Gateway Interfaces

For connectivity to traditional networks, it is necessary to connect the VNICs of the service appliances that have a gateway role assigned to them so that they can communicate with the underlay network. From vCenter, right click on the appliance and select the “Edit Settings” menu item. Then select the hardware tab, and set one network adapter to the port group of the vDS.

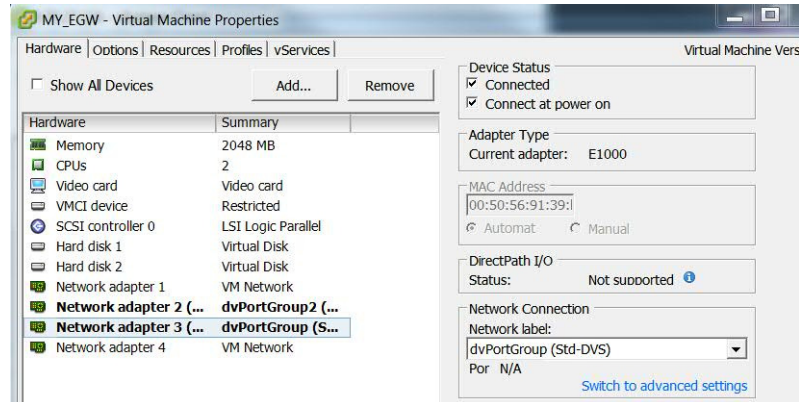


Note: Each NIC is connected to a network. For example: management network, DOVE tunnel network, and external network.

For the VLAN gateway, we then need to select another vnic and set it to attach to the VLAN that it was configured with. To do this, select the pre-configured “tagged” profile that has been configured with the proper VLANID.



For the external gateway, the second VNIC is attached via an untagged profile to a standard vDS that provides connectivity to the external network:



5000V Host Module

The 5000V vDS Host Module is deployed on ESXi hosts. It implements overlay networks support in addition to L2 switching required by VMs that wish to communicate via the SDN VE overlay networks.

Install 5000V Host Module

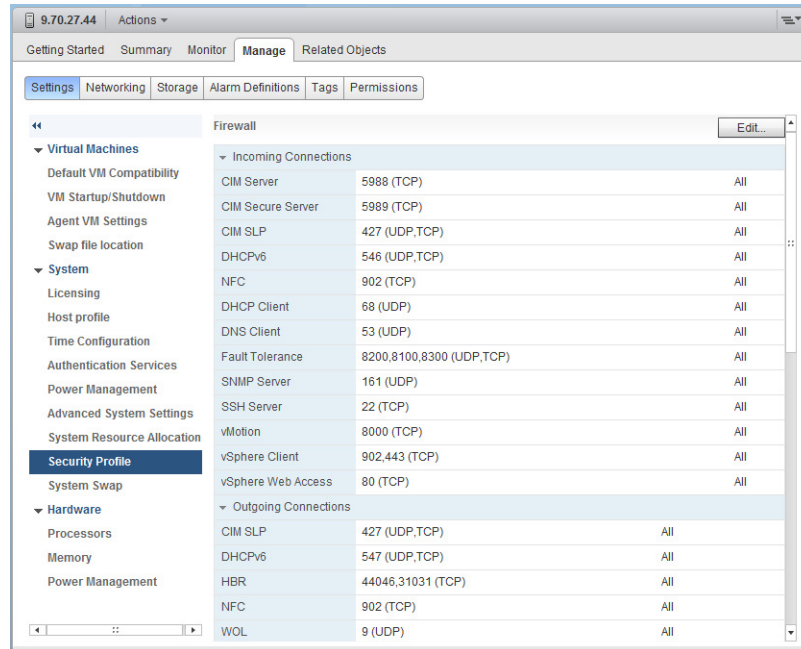
Preconditions

Verify ESXi Images

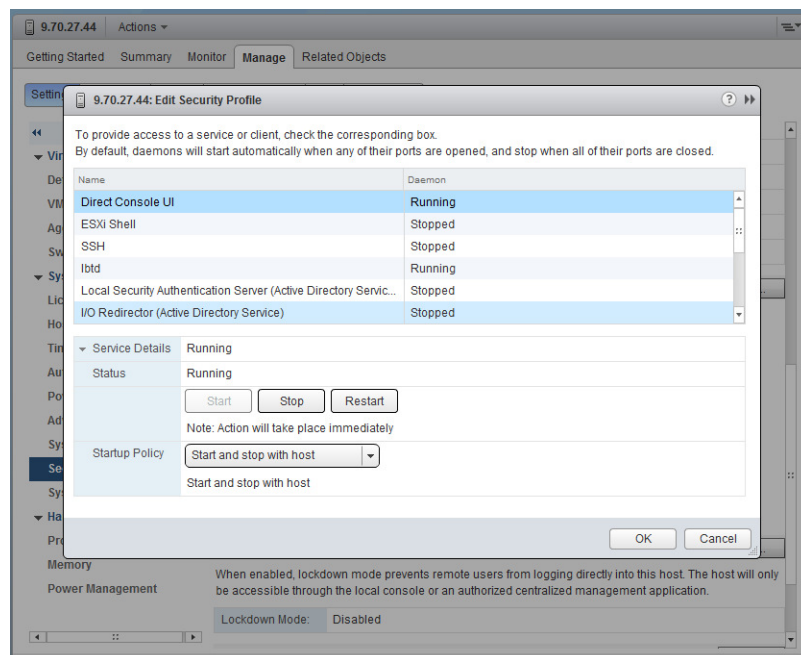
The ESXi hosts that will host the 5000V vDS Host Module must be running either VMware ESX 5.0 or 5.1 with test certificates.

Enable SSH on ESXi Hosts

To install the 5000V vDS Host Module, it is first necessary that the ESXi hosts be running ssh. To do this, go to vSphere, and select the Settings button under the Manage tab:



Scroll down to services and click the Edit button on the right.



Select ESXi Shell and SSH and start both of them, then click the OK button in the lower right. Verify by sshing to the ESXi host using the root and the ESXi root password.

Copy 5000V vDS Host Module File to ESXi Machines

Use `scp` to copy the OHM `zip` file to the destination ESXi machine:

```
[bhal@oc2072406814 Jun-03-2013]$ scp -p 5000V-host-module.zip
root@9.70.27.194:/tmp
Password:
5000V-host-module.zip
100%
```

Install 5000V vDS Host Module VIB

Use SSH to access the ESXi machine and log in as root. Then change to the directory where the zip file was copied.

Install the host module with the following command:

```
esxcli software vib install -d=file://`pwd`/5000V-host-module.zip
```

This will generate the following type of output:

```
/tmp # esxcli software vib install -d=file://`pwd`/5000V-host-module.zip
Installation Result
  Message: The update completed successfully, but the system needs to be
  rebooted for the changes to be effective.
  Reboot Required: true
  VIBs Installed: IBM bootbank ibm-esx-5000V 1.1.0-130527
  VIBs Removed: IBM_bootbank_ibm-esx-5000V 1.1.0-130513
  VIBs Skipped:
```

When this is complete, reboot the ESXi host. After it reboots, re-enable the ESXi shell and SSH service (if necessary) and make sure that the host module is installed correctly:

```
esxcli software vib list | grep -ir 5000V
```

```
/tmp # esxcli software vib list | grep -ir 5000V
ibm-esx-5000V          1.1.0-130527          IBM
VMwareAccepted 2013-5-31
```

Configure the Underlay (Physical) Networks at the Unified Controller

Underlay network is the physical network to which the uplinks of 5000V vSwitches connect. This is the network over which SDN VE encapsulated packets flow between Tunnel End Points.

Since the vSwitches do not make use of the host's IPv4 routing capabilities, it is necessary to provide the IPv4 gateway information to the vSwitches. The vSwitches can learn the IPv4 address and Subnet that was configured for the VM Kernel NIC (vmknics) that was attached to the vSwitch. However, the Default Gateway information that is present for the host cannot be used by the vSwitches since the vSwitches may be connected to a physically separate and isolated network segment from the one in which the Default Gateway exists.

For this reason, it is necessary to configure the network segments to which the TEPs will connect and the NextHop or Gateway IP that would perform IPv4 routing functions for that network segment.

Notes:

1. It is necessary to configure the network information on the Unified Controller **before** connecting vmknics to the vSwitches.
2. An underlay network configuration cannot be modified. To change the next hop, delete the net and mask combination (`underlay-network del` command) and add a new configuration.
3. Ensure that the TEPs (VMKNICs) are given addresses and netmasks that correspond to the configuration made on the Unified Controller.
4. It is not necessary to configure the underlay network unless the TEPs span multiple subnets.

The commands are as follows:

- To configure an underlay network segment:

```
SDN-VE-Controller(config)# underlay-network subnet 1.1.1.0 mask 255.255.255.0 nexthop 1.1.1.254
```

- To remove a previously configured network segment: (next hop cannot be specified)

```
SDN-VE-Controller(config)# no underlay-network id <subnet ID>
```

- To display the configured underlay networks:

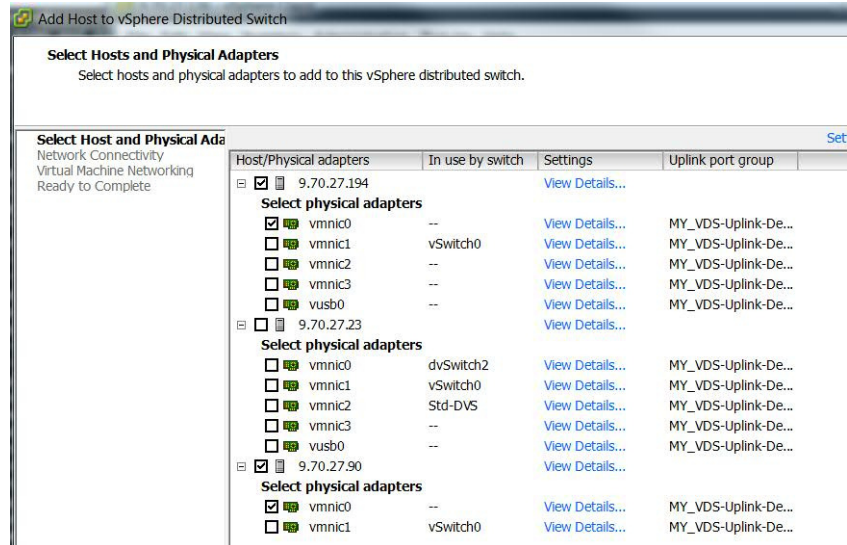
```
SDN-VE-Controller(config)# show sdnve-dove underlay-network
```

A sample output of this command:

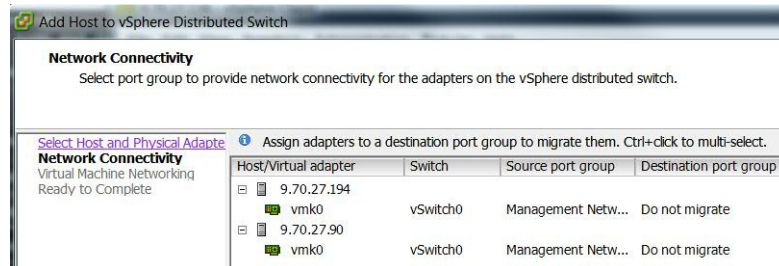
UNDERLAY NETWORK INFORMATION			
ID	IP	MASK	NEXTHOP
1	1.1.1.0	255.255.255.0	1.1.1.254

Attach ESXi Hosts to vDS

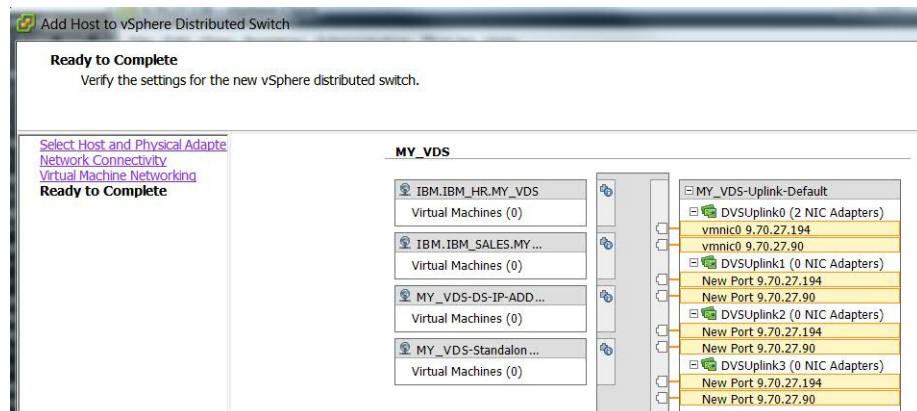
Before VMs on an ESXi host can communicate via the overlay network, the host needs to be attached to the vDS. From vCenter, go to **Home | Inventory | Networking** in the navigation bar, right click on the vDS, select “Add Host” from the menu and then select the hosts and physical adapters to add to the vDS:




Once desired hosts and physical adapters have been selected, click on “Next”:



Then click on “Next”:




Finally, click on “Finish”:

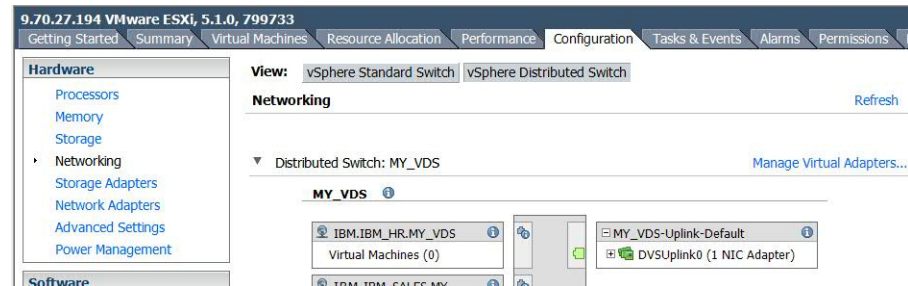


Configure TEPs


To configure the Tunnel End Point (TEP) on an ESXi host, first select **Inventory | Hosts and Clusters** in the navigation bar on vCenter, then select the host in question:



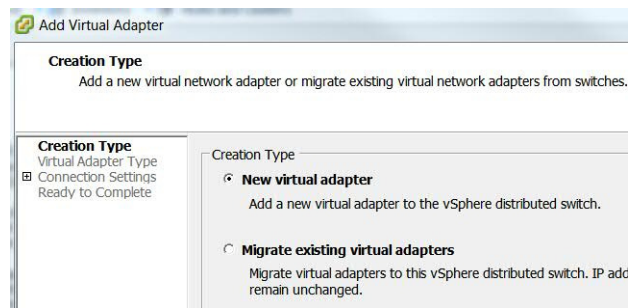
Go to the Configuration tab and select “Managed Virtual Adapters”:



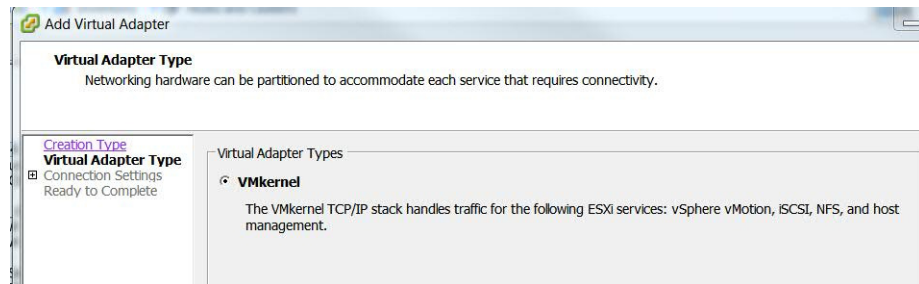
In the “Manage Virtual Adapters” window, click Add:



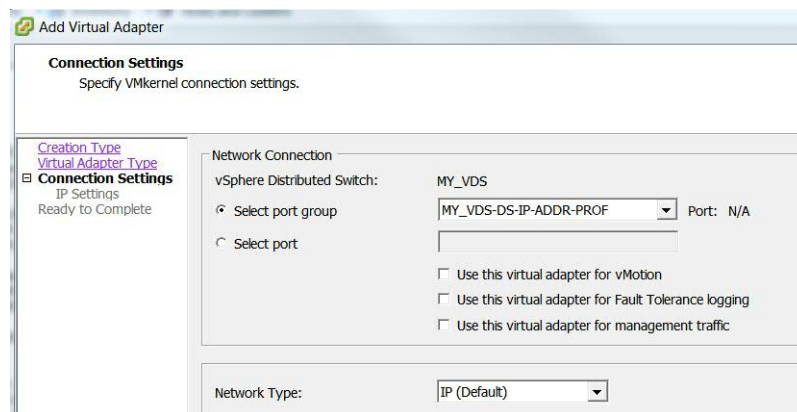
Select “New Virtual Adapter”:



Ensure that “VMkernel” is selected and click “Next”:




Then select the TEP Profile name for the Dove Tunnel as the port group and click “Next”:

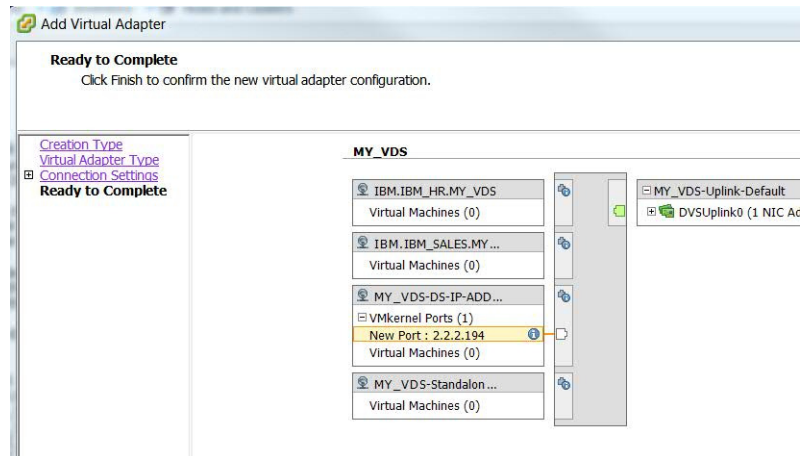


Note: Because this virtual adapter must be dedicated for TEP operation, it is required that the boxes for vMotion, Fault Tolerance logging, and management traffic must be left unchecked.


Enter the tunnel address and subnet mask and click “Next”:



Verify all the information and click Finish:



Verify that a new vmkernel adapter has been created:



At the Unified Controller, verify that the TEP is registered via the “show switch-info” command:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove switch-info
Tunnel Endpoint IP
=====
2.2.2.194
```

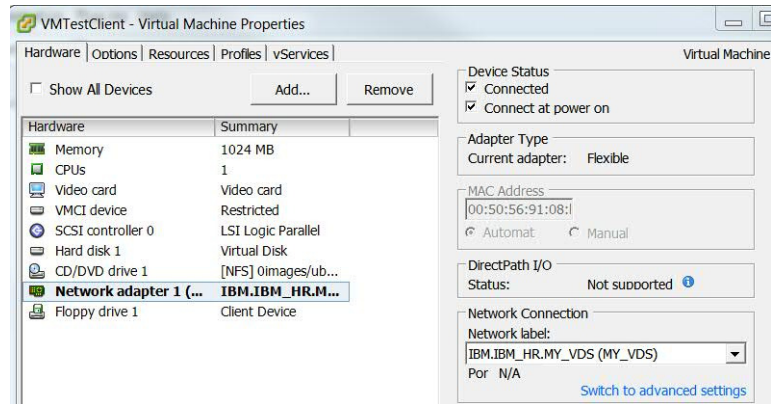
Repeat to register remaining ESXi hosts:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove switch-info
Tunnel Endpoint IP
=====
2.2.2.194
2.2.2.90
```

Attach End Systems

Note: This section assumes that end system VMs have already been deployed.

In vCenter, right click on the VM end system and select “Properties”. Change the network adapter to connect to the VDS and click “OK”:



Chapter 6. Network Services

Logical Groups

A logical group consists of ports/devices recognized by the operating system. Entities in the same logical group can exchange packets.

Subnets can also be created in logical groups.

The IBM SDN VE solution supports logical networks, subnets, and ports. The logical group service enables multi-tenant partitioning of end stations into different connectivity groups.

Multiple users with different privileges can be created for each tenant. Only administrators (system administrators and tenant administrators) can create Logical group objects (connectivity groups / subnets / ports). Each tenant has a unique ID. A DOVE administrator tenant, with an ID of 1, is created by default.

Tenant administrators can create and administer their own groups. They can move end stations from the default group into their own logical group.

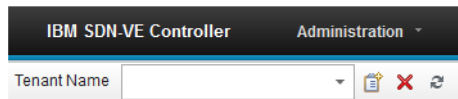
Logical groups data model is an abstract connectivity model that allows for grouping end stations from different subnets into connectivity groups. All hosts in the same connectivity group can communicate with each other; all hosts in different connectivity groups can only communicate with each other if a policy is defined between the connectivity groups.

This logical group data model can also work in an OpenStack environment with the use of Neutron APIs.

Creating a Tenant

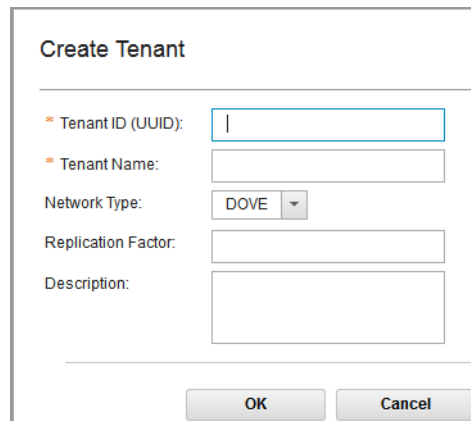
Create a tenant using the IBM Unified Controller GUI as follows:

1. Login to the controller GUI: `https://<Controller HA external IPv4 address>:8443`.
2. Select **Create** (📄) icon. The **Create Tenant** window is displayed.



The screenshot shows the 'IBM SDN-VE Controller' Administration interface. At the top, there is a dark header with 'IBM SDN-VE Controller' and 'Administration' with a dropdown arrow. Below the header, there is a form titled 'Tenant Name' with a text input field and a dropdown arrow. To the right of the input field are three icons: a document with a plus sign (Create), a red X (Delete), and a circular arrow (Refresh).

3. Specify the tenant details:

A screenshot of a 'Create Tenant' dialog box. It contains several input fields: 'Tenant ID (UUID)' with a text box containing a vertical bar, 'Tenant Name' with an empty text box, 'Network Type' with a dropdown menu showing 'DOVE', 'Replication Factor' with an empty text box, and 'Description' with a larger empty text box. At the bottom are 'OK' and 'Cancel' buttons.

Create Tenant

Tenant ID (UUID):

Tenant Name:

Network Type:

Replication Factor:

Description:

OK Cancel

Table 5. Tenant Specifications

Component	Description
Tenant ID (UUID)	Specify a unique identifier.
Tenant Name	Specify a name for the tenant.
Network Type	Select network type: DOVE or OF (Open-Flow)
Replication Factor	(For DOVE network type only) Specify the number of nodes that should replicate this information.
Description	Specify a description of the tenant.

4. Select **OK**.

Creating a Logical Group

You can create tenant-based logical groups as follows:

1. Login to the controller GUI: <https://<Controller HA external IPv4 address>:8443>.
2. Select the tenant.
3. Select **Services > Logical Groups**. Select the tenant for which you want to create the logical group.

4. Select **Create** (📄) icon. The **Create Connectivity Group** window is displayed.

Create Connectivity Group

ID:

VNID:

Group Name:

Admin State Up:

true

Group Type:

dedicated

isNeutron:

false

Traffic Type:

BEST_EFFORT

Precedence Type:

FLASH_OVERRIDE

Limit Delay:

Limit Throughput:

Limit Reliability:

Average Rate:

Peak Rate:

Burst Rate:

OK

Cancel

Specify the connectivity group details:

Table 6. Connectivity Group Specifications

Component	Description
ID	UUID of the group (1-36 alphanumeric characters)
VNID	Virtual Network ID for the group; Unsigned integer; Specify in range 1-65535
Group Name	Specify a name for the group.
Admin State Up	The administrative state of group. If false (down), the group does not forward packets.
Group Type	<p>Specify if the group resource is dedicated/shared/external.</p> <p>Dedicated: The group is dedicated to the tenant, but can be reused for another tenant.</p> <p>Shared: The group is shared with the underlay network and cannot be reused anywhere.</p> <p>External: External groups can communicate with external networks.</p>
isNeutron	Specify if the group uses OpenStack Neutron APIs.

Table 6. Connectivity Group Specifications

Component	Description
Traffic Type	Select the traffic type for the group: BEST_EFFORT BACKGROUND EXCELLENT_EFFORT CRITICAL_APPLICATIONS VIDEO VOICE INTERNETWORK_CONTROL NETWORK_CONTROL
Precedence Type	Select the type of precedence: ROUTINE PRIORITY IMMEDIATE FLASH FLASH_OVERRIDE CRITIC_ECP INTERNETWORK_CONTROL NETWORK_CONTROL
Limit Delay	Accepts an integer value: 0 or 1. 0 - False 1 - True
Limit Throughput	Accepts an integer value: 0 or 1. 0 - False 1 - True
Limit Reliability	Accepts an integer value: 0 or 1. 0 - False 1 - True
Average Rate	The average number of kilobytes per second (KBps) to allow across a port or a portgroup.
Peak Rate	The number of kilobytes per second (KBps) to allow across a port or a portgroup, when it is sending/receiving a burst of traffic.
Burst Rate	Maximum number of kilobytes to allow in a burst.

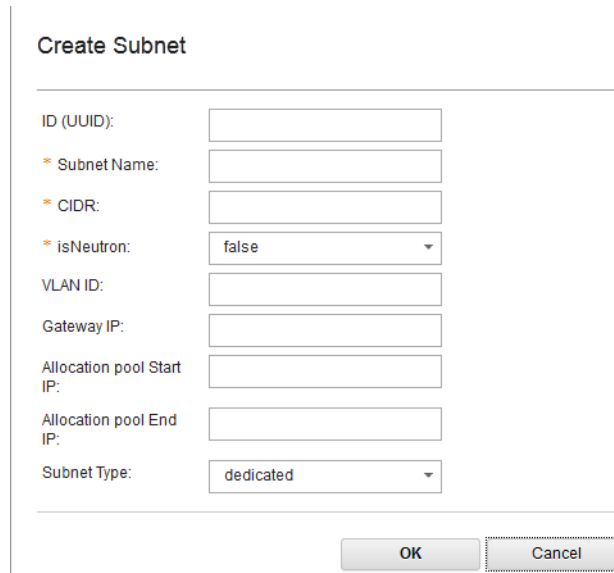
5. Select **OK**.

Creating a Subnet

A subnet represents an IP address block that can be used to assign IP addresses to virtual instances. You can create a subnet in a logical group.

1. Login to the controller GUI: <https://<Controller HA external IPv4 address>:8443>.
2. Select the tenant.
3. Select **Services > Logical Groups**.
4. Select the tenant for which you want to create the subnet.
5. Expand *<tenant name>*.
6. Select **Subnets**. You will see the **List of Subnets** screen in the right pane.

7. Select **Create** (🔑) icon. The **Create Subnet** window is displayed.



Specify the subnet details:

Table 7. Subnet Specifications

Component	Description
ID (UUID)	UUID of the subnet (1-36 alphanumeric characters)
Subnet Name	Specify a name for the subnet
CIDR	Specify the IPv4 address for the subnet and the routing prefix. Example: 9.110.20.32/24
isNeutron	Specify if the subnet uses OpenStack Neutron APIs.
Gateway IP	Specify the gateway IP address for the subnet.
Allocation Pool Start IP	Specify the first IP address of the range of IP addresses allocated to the subnet.
Allocation Pool End IP	Specify the last IP address of the range of IP addresses allocated to the subnet.
Subnet Type	Select if the subnet will be dedicated to a connectivity group or shared between connectivity groups.

Create a Port

Note: Only Layer 3 ports can be created for DOVE tenants.

A port represents a virtual switch port on a logical group switch. Virtual instances attach their interfaces to ports. The logical port also defines the MAC address and the IP addresses that you must assign to the interfaces that are plugged into the port. You can create an Layer 2 or Layer 3 port in a logical group.

You can create ports as follows:

1. Login to the controller GUI: <https://<Controller HA external IPv4 address>:8443>.
2. Select the tenant.
3. Select **Services > Logical Groups**. Select the tenant for which you want to create the ports.
4. Expand <tenant name>. You will see the list of connectivity groups, if any.
5. Select <connectivity group name>. The **List of Ports belong to Group: <connectivity group name>** page is displayed in the right pane.
6. Select **Create** (📄) icon. The **Create L2/L3 Port** window is displayed.

Create L2/L3 Port

ID (UUID):

* Name:

* MAC:

Admin State Up:

IP Address:

Specify the port details:

Table 8. Port Specifications

Component	Description
ID (UUID)	UUID of the port (1-36 alphanumeric characters)
Name	Specify a name for the port
MAC	Specify the MAC address.
Admin State Up	Select if the port should be up (<code>true</code>) or down (<code>false</code>).
IP Address	Specify an IP address for the port.

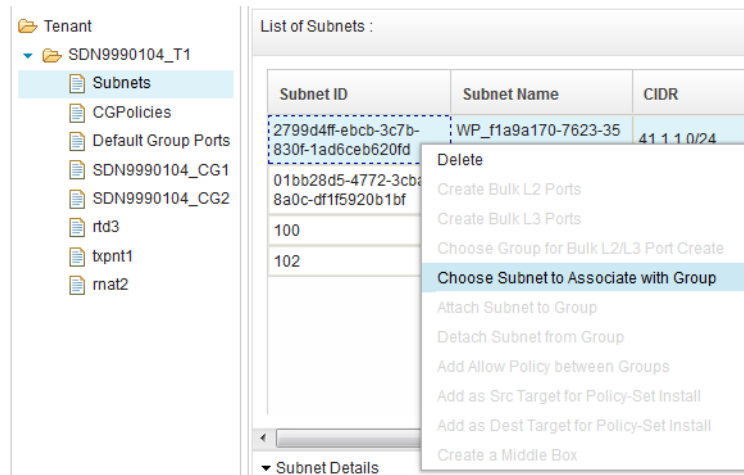
7. Select **OK**.

Assign Subnet to a Connectivity Group

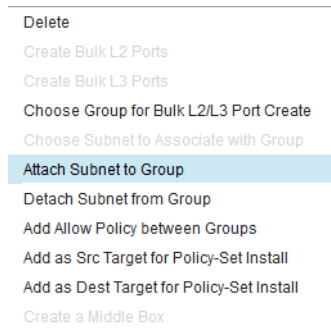
You can assign a subnet to a connectivity group as follows:

1. Login to the controller GUI: <https://<Controller HA external IPv4 address>:8443>.
2. Select the tenant.
3. Select **Services > Logical Groups**. Select the tenant.
4. Expand <tenant name>. You will see the list of connectivity groups in the right pane.
5. Select a connectivity group in the right pane.
6. Select **Subnets** in the left pane. The list of subnets is displayed in the right pane.

7. Right-click on the subnet in the right pane.



8. Select **Choose Subnet to Associate With Group**.
9. Select **Services > Logical Groups > <tenant name>**. The connectivity groups are displayed on the right pane.
10. Right-click on the connectivity group and select "Attach Subnet to Group".




Layer 3 Configuration

Note: This section is applicable only to OpenFlow tenants.

When you create an OpenFlow tenant, a logical router is automatically created.

You can view the list of routers as follows:

1. Login to the controller GUI: <https://<Controller HA external IPv4 address>:8443>.
2. Select the tenant.
3. Select **Services > L3**. The routers are displayed in the right pane.



For each subnet that you create, a virtual interface is automatically created.

Name:	CG_2
Admin Status:	UP
Default Gateway:	
V Mac:	00:50:56:a2:40:71

Virtual Interfaces		Routing Table				
Interface ID	Interface Name	Admin Status	VirtualMac	IP Address	Subnet Mask	Vlan
5b914b6a-b5af-44a5-a3e5-b302bcae42ae	CG_9f9cbb8f-b5db-4be3-be6f-cc89f8bd4bc4	UP	00:50:56:a2:40:71	1.1.1.1	255.255.255.0	Untagged
9630fcc3-ca16-4f34-b152-0d5e0074bd06	CG_2	UP	00:50:56:a2:40:71	2.2.2.1	255.255.255.0	Untagged

You can add a route as follows:

1. Select the Routing Details tab.
2. Select **Create** (📄) icon. The **Create Route** window is displayed.

Create Route

* Network Address:

4.4.4.10

* Subnet Mask:

255.255.255.0

* NextHop/Gateway:

4.4.4.1

OK

Cancel

3. Specify the route details.
4. Select **OK**.

Connectivity Group Policy

Adding a Policy Between Two Connectivity Groups

You can configure a policy between two connectivity groups. The communication could be unidirectional or bidirectional.

1. Login to the controller GUI: <https://<Controller HA external IPv4 address>:8443>.
2. Select the tenant.
3. Select **Services > Logical Groups**.
4. Select *<tenant name>*. The connectivity groups are displayed in the right pane.
5. Select the first connectivity group.
6. Press and hold **ctrl** and select the second connectivity group.
7. Right-click on the selection.
8. Select Add Allow Policy between Groups. The **Add Connectivity Group Policy** window is displayed.

9. Select the communication: UNI_DIRECTIONAL or BI_DIRECTIONAL.

Add Connectivity Group Policy

Group 1 ID:

4

Group 2 ID:

5

Traffic Type:

UNICAST

Directional:

UNI_DIRECTIONAL

SNAT Start IP:

SNAT End IP:

SNAT Start Port:

SNAT End Port:

OK

Cancel

Specify the SNAT Pool details.

Table 9. SNAT Pool Specification

Component	Description
SNAT Start IP	Starting IP address of the range of addresses you want to allocate for NAT.
SNAT End IP	Ending IP address of the range of addresses you want to allocate for NAT.
SNAT Start Pool	Starting port number of the ports to be assigned for NAT. Port numbers can be in the range: 1-65535
SNAT End Pool	Starting port number of the ports to be assigned for NAT.

10. Select **OK**.

Monitor/Redirect Sessions

Flow Replication and Redirection is a monitoring and troubleshooting service. It helps in setting up replication/redirection path source, destination, and replication/redirection destination in the network. Replication is similar to Switched Port Analyzer (SPAN) functionality available in network switches.

You can specify a flow based on any combination of source IP/MAC address, destination IP/MAC address, and transport protocol. You should specify the replication/redirection destination, and an end station to which the replicated/redirected flow needs to be sent. Flow Replication service determines a replication point—a point from which a flow is replicated and sent to the replication destination.

If there are any network changes in the path from replication point to replication destination, the replication service automatically switches to a new path and reprograms the switches.

Create a Replication/Redirection Session

You can create a replication/redirection session as follows:

1. Login to the controller GUI: `https://<Controller HA external IPv4 address>:8443`.
2. Select the tenant.
3. Select **Services > Monitor/Redirect Sessions**. The session list is displayed in the right pane.
4. Select **Create** (📄) icon. The **Create Session Replication/Redirection** window is displayed.

Create Session Replication/Redirection

* Session Name:

* Session Mode: Redirect ▾

* Source/Destination Tenant:

Protocol: ANY ▾

Source: IP ▾ Destination: IP ▾ Target: IP ▾

IP: IP: IP:

Port: Port: Tenant:

OK Cancel

Specify the session details:

Table 10. Session Specifications

Component	Description
Session Name	Specify a session name.
Session Mode	Select session type: Redirect or Replicate.
Source/Destination Tenant	Default value is displayed. This cannot be edited.
Protocol	Select protocol: Any, ICMP, TCP, UDP
Source	Select source type: IP, MAC Specify the IP address or MAC address, as applicable.
Destination	Select destination type: IP, MAC Specify the IP address or MAC address, as applicable.
Target	Select target type: IP, MAC Specify the IP address or MAC address, as applicable. Specify the tenant name or ID.

5. Select **OK**.

Start/Stop/Delete a Monitor/Redirect Session

You can start, stop, or delete a session as follows:


1. Login to the Controller GUI: `https://<Controller HA external IPv4 address>:8443`.
2. Select the tenant.
3. Select **Services > Monitor/Redirect Sessions**. The list of sessions is displayed in the right pane.
4. Right-click on the session.
5. Select **Start**, **Stop**, or **Delete**, as required.

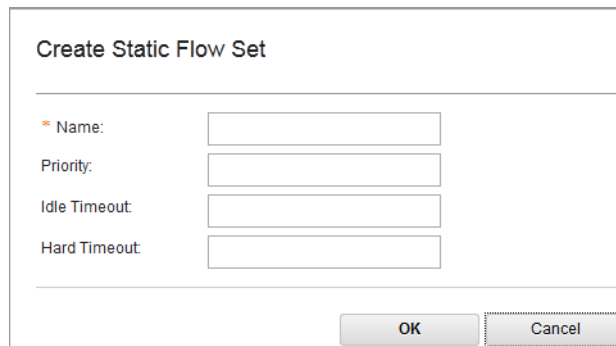
Static Flows

You can create a Static Flow Group, which in turn creates one or more Static Flow(s) within the Static Flow Group. You can install, uninstall, or delete Static Flow Groups, as required.

Create a Static Flow Set

You can create a static flow set as follows:

1. Login to the Controller GUI: `https://<Controller HA external IPv4 address>:8443`.
2. Select a tenant.
3. Select **Services > Static Flows**. The **Flow Set Details** window opens in the right pane.
4. Select the Create () icon. The **Create Static Flow Set** window is displayed.



The image shows a 'Create Static Flow Set' dialog box. It has a title bar with the text 'Create Static Flow Set'. Below the title bar, there are four labeled text input fields: 'Name:', 'Priority:', 'Idle Timeout:', and 'Hard Timeout:'. At the bottom right of the dialog box, there are two buttons: 'OK' and 'Cancel'.

Specify the following:


Table 11. Static Flow Details

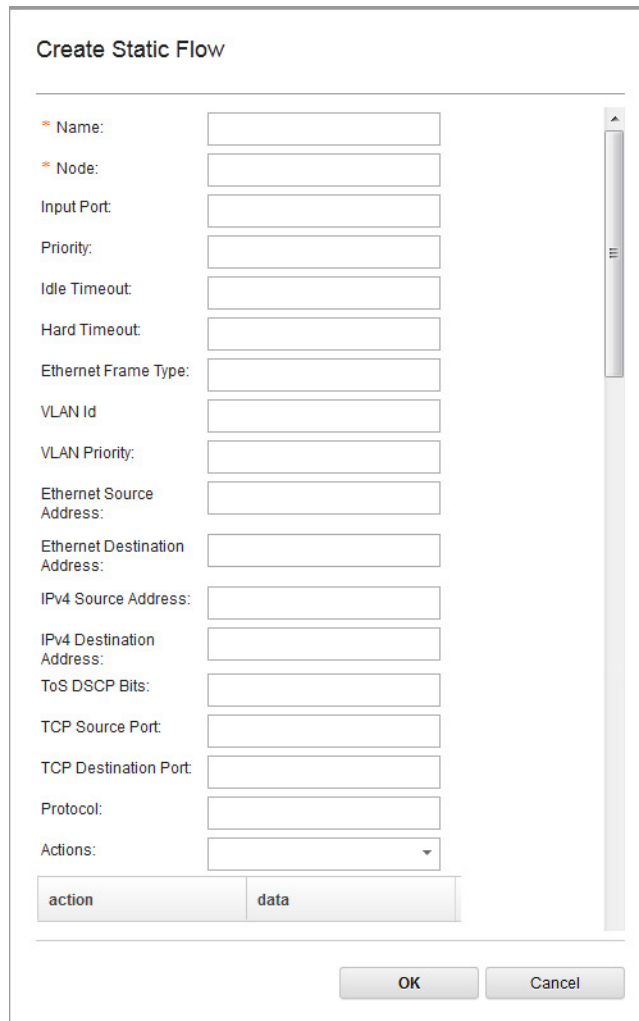
Component	Description
Name	Unique name for the static flow.
Priority	Specify the priority for this set.
Idle Timeout	Specify the idle timeout value.
Hard Timeout	Specify the hard timeout value.

5. Select **OK**.

Create a Static Flow

You can create a static flow as follows:

1. Login to the Controller GUI: <https://<Controller HA external IPv4 address>:8443>.
2. Select a tenant.
3. Select **Services > Static Flows**. The **Flow Set Details** window opens in the right pane.
4. Select the flow set to which you want to add a static flow.
5. Select the Create () icon from the top-right corner of the **Static Flow List** display area. The **Create Static Flow window** is displayed.



The 'Create Static Flow' window is a modal dialog box with a title bar. It contains a list of input fields for configuring a static flow. The fields are: Name (required), Node (required), Input Port, Priority, Idle Timeout, Hard Timeout, Ethernet Frame Type, VLAN Id, VLAN Priority, Ethernet Source Address, Ethernet Destination Address, IPv4 Source Address, IPv4 Destination Address, ToS DSCP Bits, TCP Source Port, TCP Destination Port, Protocol, and Actions (a dropdown menu). Below the input fields are two tabs: 'action' and 'data'. At the bottom right are 'OK' and 'Cancel' buttons.

Specify the following:

Table 12. Static Flow Details

Component	Description
Name	Unique name for the static flow.
Node	Switch ID. Format: xx:xx:xx:xx:xx:xx:xx:xx:xx

Table 12. Static Flow Details

Component	Description
Input Port	Ingress port for the flow.
Priority	Priority value of the flow set (0-65535 seconds)
Idle Timeout	Idle timeout value (0-65535 seconds). If no match is found for the flow for the configured time, the flow is removed from the table.
Hard Timeout	Hard timeout value (0-65535 seconds). Flow is removed from the table after the configured time irrespective of the match status.
Ethernet Frame Type	Ethernet type for the flow.
Vlan ID	Out VLAN ID for the flow. 0 to be used for untagged packets. Values: 0-4094
Vlan Priority	VLAN priority for the flow. Values: 0-7
Ethernet Source Address	Source MAC address for the flow.
Ethernet Destination Address	Destination MAC address for the flow.
IPv4 Source Address	Source IP address for the flow.
IPv4 Destination Address	Destination IP address for the flow.
ToS DSCP Bits	Type of service for the flow. 0 - lowest priority; 63 - highest priority.
TCP Source Port	Ingress port for the flow.
TCP Destination Port	Egress port for the flow.
Protocol	IP protocol for the flow. Values: 0-255
Actions	<p>Select the action for the flow:</p> <p>DROP Drop the flow.</p> <p>OUTPUT Out port(s)</p> <p>STRIP_VLAN Out VLAN ID (0-4094) 0 to be used for untagged packets</p> <p>SET_DL_SRC Source MAC address</p> <p>SET_DL_DST Destination MAC address</p> <p>SET_VLAN_ID VLAN ID</p> <p>SE_VLAN_PCP VLAN priority (0-7)</p> <p>SET_DL_TYPE Ethernet type</p> <p>SET_NW_SRC Network source (IP address)</p> <p>SET_NW_DST Network destination (IP address)</p> <p>SET_NW_TOS Type of service (0-63) 0 - lowest priority; 63 - highest priority</p> <p>SET_TP_SRC Transport layer source</p> <p>SET_TP_DST Transport layer destination.</p>

6. Select **OK**.

Install a Static Flow Set

Install a static flow set as follows:

1. Login to the Controller GUI: <https://<Controller HA external IPv4 address>:8443>.
2. Select a tenant.
3. Select **Services > Static Flows**. The **Flow Set Details** window opens in the right pane.
4. Right-click on the flow set you want to install.
5. Select **Install**.

Delete a Static Flow or Uninstall a Flow Set

Delete a static flow as follows:

1. Login to the Controller GUI: <https://<Controller HA external IPv4 address>:8443>.
2. Select a tenant.
3. Select **Services > Static Flows**. The **Flow Set Details** window opens in the right pane. The static flows are displayed in the **Static Flow List** section.
4. Right-click on the flow or flow set you want to remove.
5. Select **Delete** or **Uninstall**.













Chapter 7. Topology

You can use the Topology Manager to view a map of the topology and the interconnection of switches and hosts in the logical groups and physical networks. You can also use the Topology Manager to export the discovered topology information.

You can make use of REST API for import the topology from external file.

Topology Manager

You can view the topology and the interconnection of the different components within the logical group or the physical network in the Topology tab. The following table describes some of the operations with the icon or labels and descriptions of the operations.

Icon or Label	Description
Refresh	Refreshes the topology view
	Saves the customizations
	Prints the topology view
	Selects part of or the complete topology view
	Zooms in the selected part of the topology view
	Zooms in the topology view
	Zooms out the topology view
	Fits the content in the window
	Displays the topology view in tree format
	Displays the topology view in hierarchical format
	Displays the topology view in force-directed format
	Displays the topology view in short-link format
	Displays the topology view in long-link format

Actions

You can use the Actions menu item in the Topology tab to perform all the actions that are represented as icons in the topology view. You can also use this menu item to export the discovered topology information and save the topology information.

To export the topology information as an image in a .jpg file or a .png file, select **Actions > Export** as Image and save the image.

To export the topology information as a .csv file, select **Action > Export** and save the file.

Search

You can use the Search menu item in the Topology tab to search for various components in the network. The following table describes the options that you can use when you click the Search menu item.

Option	Description
Filter	Type the name of the component that you want to search in this field
Center in View	Click to center the selected component in the Topology Manager
Highlight	Click to highlight the selected component in the Topology Manager
Clear Search	Click to clear the filters in the search window and display the default view in the Topology tab
Close	Click to close the Search window

Logical Groups

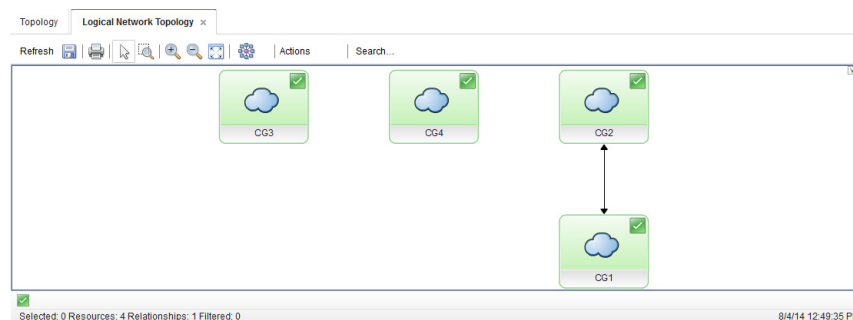
Logical groups provide multi-tenancy service. You can view the logical group that is specific to the selected tenant in the IBM Unified Controller GUI.

Viewing the Logical Group

You can view the logical group and the components specific to the tenant such as the subnets, ports, and routers.

View the logical group and the components as follows:

1. Select the tenant.
2. Select **Topology > Logical Network**. The logical groups and the components are displayed in the right pane.



Viewing Logical Group Properties

You can view properties of the logical groups and the components such as the subnets and ports.

To view the properties of the logical group and the components, complete the following steps:

1. Select the tenant.
2. Select **Topology > Logical Network**. The logical network and the components are displayed in the right pane.
3. Right-click on the logical group or the component and select **Properties**.

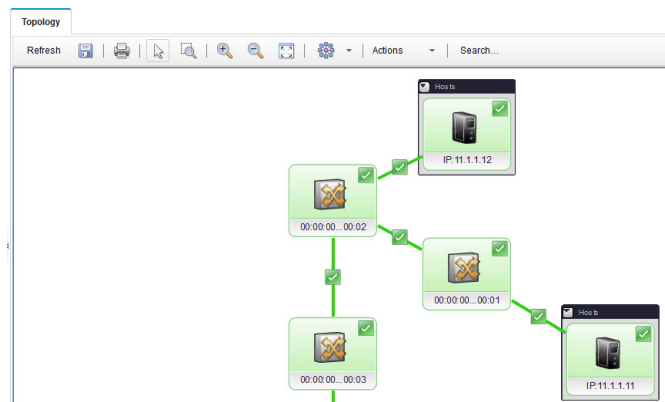
Physical Networks

You can view the physical network in the IBM Unified Controller GUI. The physical network is not specific to any tenant.

Viewing the Physical Network

You can view physical network components, such as switches and hosts, as follows:

Select **Topology > Physical Network**. The physical network and the components are displayed in the right pane.

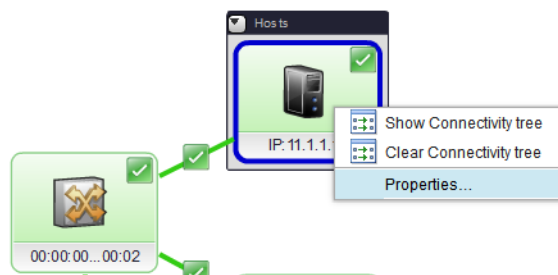


Viewing Properties of the Physical Network

You can view properties of the components, such as the switches and hosts, in the physical network.

On the controller GUI:

1. Select **Topology > Physical Network**. The physical network and the components are displayed in the right pane.
2. Right-click on the component and select **Properties**.

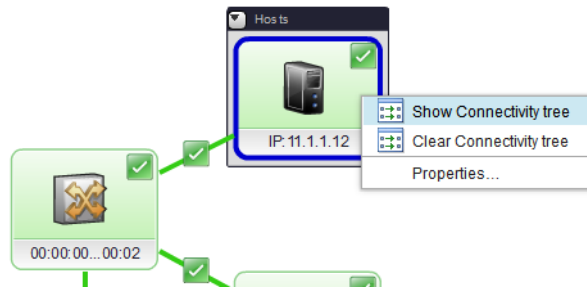


Viewing the Connectivity Tree

You can view the connectivity between the selected host and the switches in the physical network.

On the controller GUI:

1. Select **Topology > Physical Network**. The physical network and the components are displayed in the right pane.
2. Right-click on the host and select the **Show Connectivity tree**. The connectivity tree is displayed.



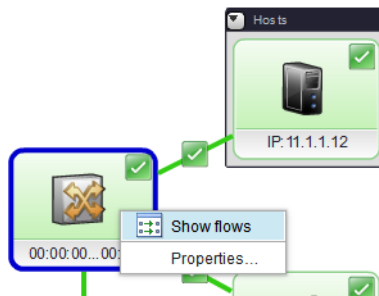
To clear the connectivity details, right-click on the host and select the **Clear Connectivity tree**.

Viewing the Flows for a Switch

You can view the flows for a switch in the physical network.

On the controller GUI:

1. Select **Topology > Physical Network**. The physical network and the components are displayed in the right pane.
2. Right-click on the switch and select the **Show flows**. The flow list for the switch is displayed.



Chapter 8. System Administration

The user with “System-Admin” role can only perform administrative tasks. Hence, only the users with this role should be able to access “Administration” (and submenus) from GUI.

To elaborate,

- Configure system settings such as license information and LDAP configuration.
- Monitor and debug logs
- Define and manage users

User Management

You can define users and roles in the IBM Unified Controller GUI. You can define the following type of users:

User type	Description
System-admin	The super-user. This user has complete access to all configuration and exec commands.
Tenant-admin	Admin for the tenant. This user can perform administrative actions for the Tenant network. Tenant-admin users can configure and view the tenant networks, subnets, ports, routers, or policies.
Tenant-operator	Tenant operator can view tenant-specific networks, subnets, ports, routers, or policies.

Note: The user with System-Admin privilege only can perform administrative tasks. Therefore, only the users with this role can access the Administration menu (and submenus) from the GUI.


Creating users

To create users, complete the following steps:

On the navigation pane of the IBM Unified Controller GUI:

1. Select **Administration > Security > User management**.

The User List window appears.



2. Select the Create User icon  .

3. Enter the user details on the Create User page:

Create User

* User Name:

Password:

Confirm Password:

* Role:

Tenant-Operator ▼

* Tenant ID:

▼

Tenant Name:

OK

Cancel

Note:

- User name: specify a user name of maximum 8 characters length. Should be alphanumeric with no special characters.
- Password: specify a password without any special characters, and confirm it.
- Role: the type of user.
- Tenant ID: the tenant to which the user will have access to.

Note: The tenant ID is not required for `System-Admin` role.

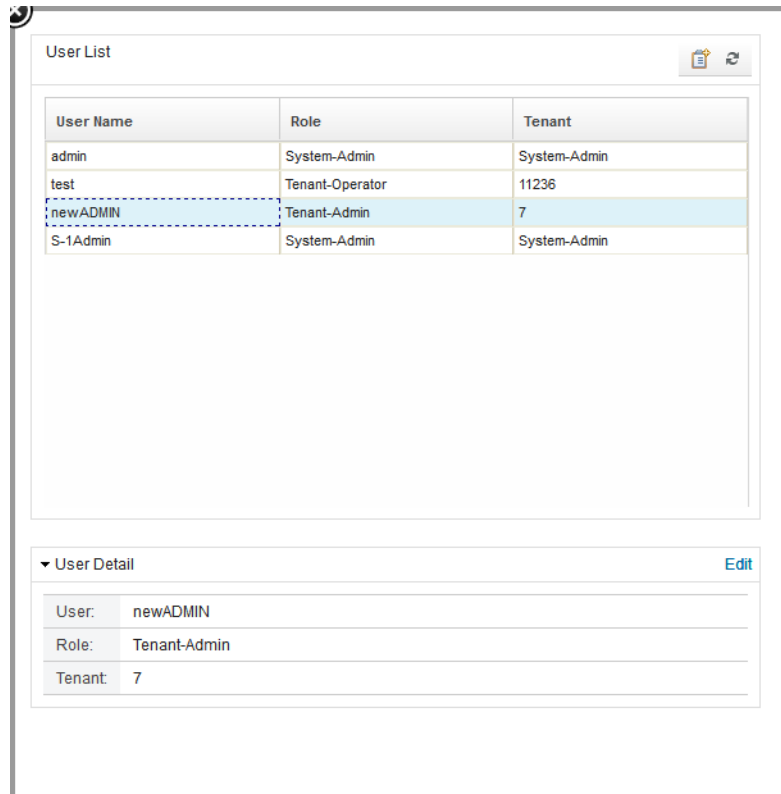
Editing users

To edit an existing user, complete the following steps:

On the navigation pane of the IBM Unified Controller GUI:

1. Select **Administration > Security > User Management**.

In the User List window that appears, select a user to display the user details in the User Detail section.



User List

User Name	Role	Tenant
admin	System-Admin	System-Admin
test	Tenant-Operator	11236
newADMIN	Tenant-Admin	7
S-1Admin	System-Admin	System-Admin

▼ User Detail [Edit](#)

User:	newADMIN
Role:	Tenant-Admin
Tenant:	7

2. Select the **Edit** button in the User Detail section to change the user details.

The screenshot shows the 'User List' window with a table of users. The 'newADMIN' user is selected. Below the table is the 'User Detail' section, which is currently collapsed. The 'User Detail' section contains fields for 'User', 'Role', and 'Tenant', all of which are populated with the details of the selected user.

User Name	Role	Tenant
admin	System-Admin	System-Admin
test	Tenant-Operator	11236
newADMIN	Tenant-Admin	7
S-1Admin	System-Admin	System-Admin

▼ User Detail

User: newADMIN

Role: Tenant-Admin

Tenant: 7

Save Cancel


3. Select **Save** to save the changes.

Deleting users

To delete an existing user, complete the following steps:

On the navigation pane of the IBM Unified Controller GUI:

1. Select **Administration > Security > User Management**.
2. In the User List window, right-click on a user and select **Delete**.



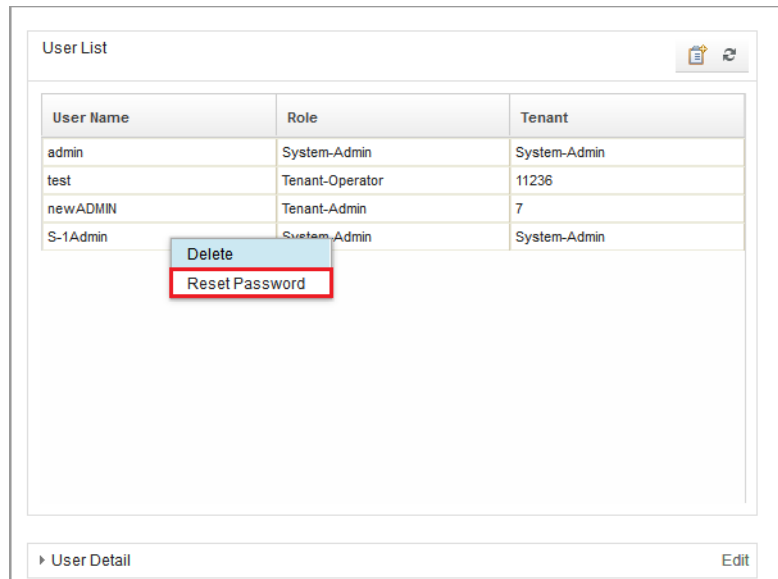
The user is deleted.

Resetting Passwords

To reset the password of an existing user, complete the following steps:

On the navigation pane of the IBM Unified Controller GUI:

1. Select **Administration > Security > User management**.
2. In the User List window, right-click on a user and select **Reset Password**.




The password for the selected user is reset to the default password (such as "welcome").

Changing Passwords

After logging in, the user can change the log on password as follows:

On the navigation pane of the IBM Unified Controller GUI:

1. Select **Admin > Change Password**.



2. Enter the new password, confirm it, and click **OK** to save changes.


The screenshot shows the 'Change Password' dialog box. It contains three input fields for password entry, each with a red asterisk icon to its left:

- Current: [password field]
- New: [password field]
- Confirm new: [password field]

At the bottom of the dialog, there are two buttons: 'OK' and 'Cancel'.

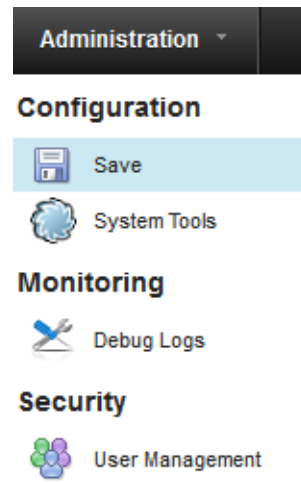
Logging out of IBM Unified Controller GUI

To log out of IBM Unified Controller GUI, select **Admin > Logout** on the navigation pane of the IBM Unified Controller GUI.



Save Configuration

To save the application configuration into disk files, select **Administration > Configuration > Save** on the IBM Unified Controller GUI.



If the configuration is successfully saved, a “Configuration Saved” message box is displayed.

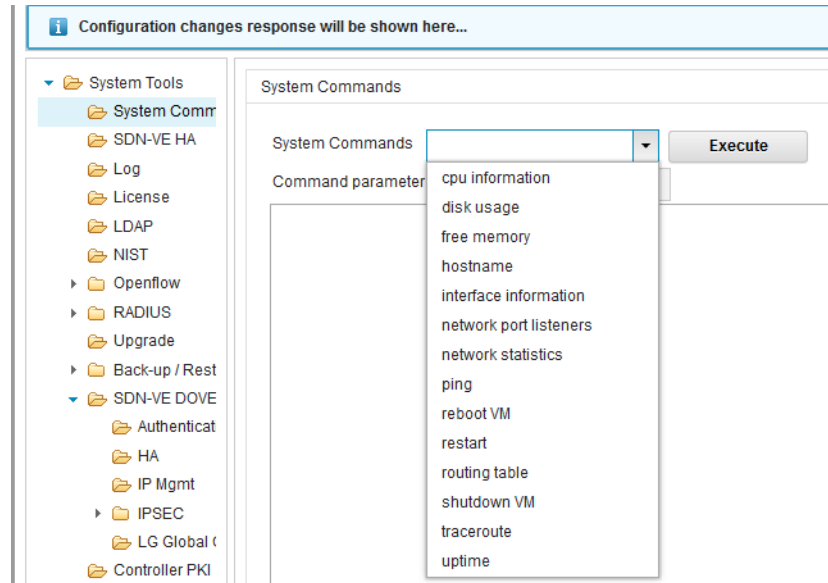


System Commands

This section lists the supported system commands.

1. On the IBM Unified Controller GUI, select **Administration > System Tools**.
2. Select **System Tools > System Commands**.

3. Select the **System Commands** drop down menu on the right pane to view the available commands.



The system commands and the command parameter are as follows.

System Command	Command Parameter	Output Description
CPU Information	Nil	CPU details such as Vendor, Family, and Model Name for each processor.
Disk Usage	Nil	Mount Point used, 1k Blocks, and Availability details
Free Memory	Nil	Total and free memory
Host Name	Nil	Name of the host
Interface Information	Nil	For each interface (loopback & etehrnet), provides info like IP Address, Netmask, Gateway, and Rx / Tx packets / errors etc
Network Port Listeners	Nil	Port, Local Address, and Remote
Network Statistics	Nil	IP, ICMP, TCP, and UDP statistics
Ping	Target IP Address	Ping Statistics
Reboot VM	Nil	Upon confirmation, it reboots the VM.
Restart	Nil	Restart the processes affected by a configuration change.
Routing Table	Nil	Destination Type, Device, and Gateway
Shutdown VM	Nil	Upon confirmation, it shuts down the VM.
Trace Route	Target IP Address	Trace route details
Uptime	Nil	System Uptime details

SDN VE HA Cluster Management

For adding a cluster or viewing cluster information, see [“Establish SDN VE Controller HA” on page 43](#).

This section provides additional details on cluster management.

Rejoining a Cluster

If a controller gets disconnected from the cluster and you need to add it back to the cluster, perform the following steps:

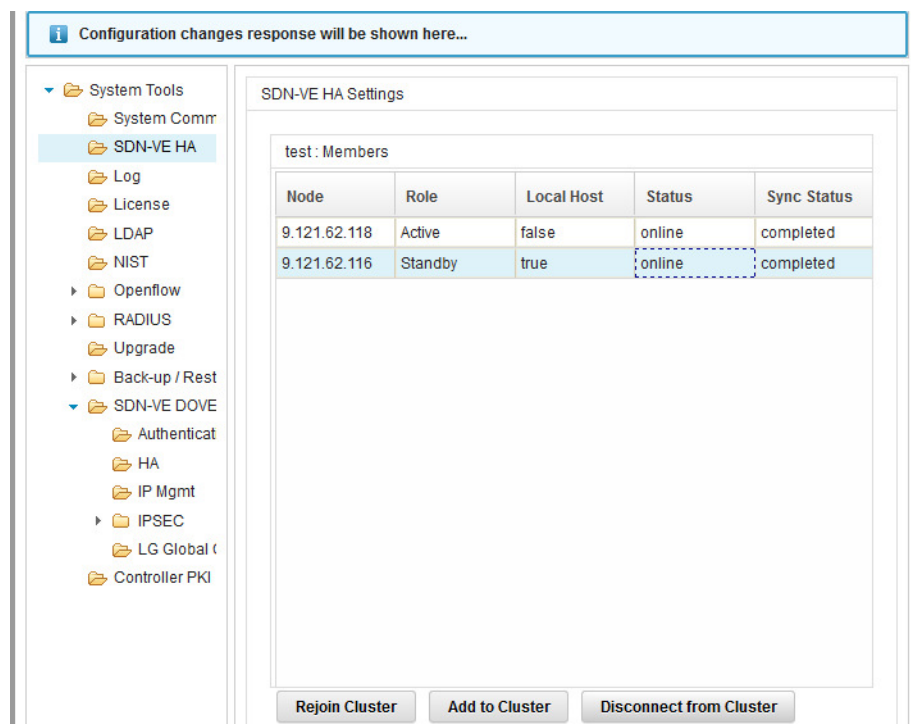
1. Add the cluster by following steps 1-2 of the section [“Establish SDN VE Controller HA” on page 43](#).
2. Select **Rejoin Cluster**.

If successful, a “Cluster rejoin completed” message is displayed on the top left corner.

Disconnecting from a Cluster

You can disconnect a specific controller from the cluster as follows:

1. Select the controller on the **SDN-VE HA Settings** display window.
2. Select **Disconnect from Cluster**.



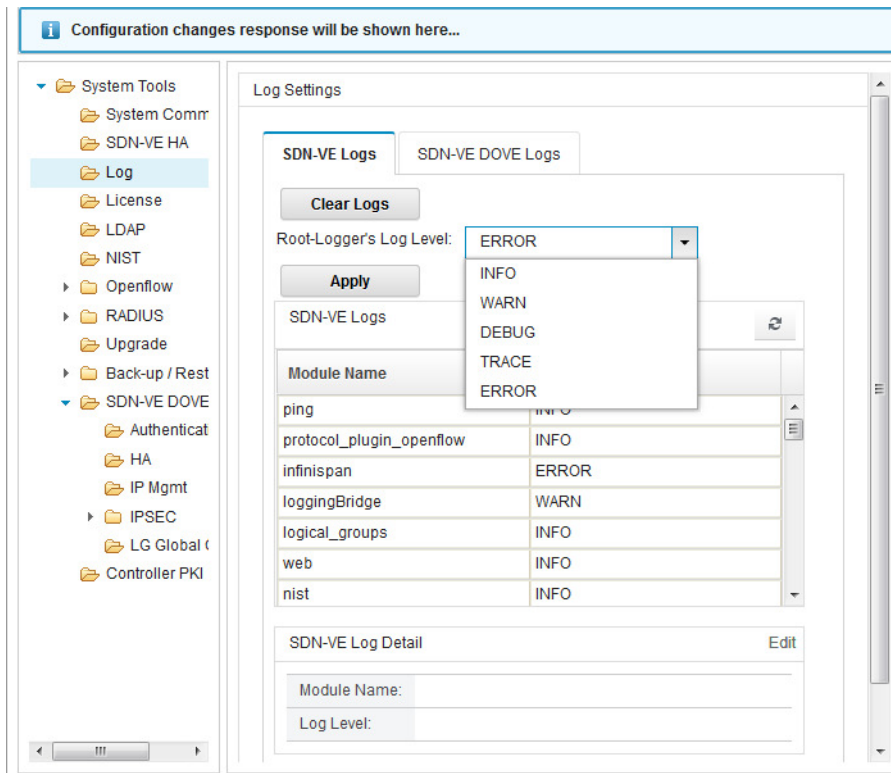
If the controller is successfully disconnected, a message is displayed on top left corner: The Node has been disconnected from the cluster - *<clusterName>*.

Log Settings

This section provides information on changing the log settings. The default log level setting is **ERROR**.

On the IBM Unified Controller GUI:

1. Select **Administration > System Tools > Log**.
2. Open the log level drop-down menu on the right pane.



3. Select the desired level and then select **Apply**.

Remote Server Setup (LDAP / RADIUS)

The remote server must be configured with the following settings:

User name (Datatype: String)	Password (Datatype: String)	Tenant ID (Datatype: String)	Role (Datatype: String)
LDAP			
alphanumeric	Type: only "simple" (LDAP mode) plain text password supported	Where to get the data from? Tenant ID should be the ID of the Tenant created. Which attribute of the LDAP user query string is used? "ou" should have the TenantID.	Where to get the data from?. The groupofNames LDAP entity should have the commonName as "System-Admin", "Tenant-Admin", "Tenant-Operator " Which attribute of the LDAP user query string is used to identify the role? Within each groupOfNames identi-fied the given user will be searched. Given user must not repeat, if it does the behavior is not consistent.

User name (Datatype: String)	Password (Datatype: String)	Tenant ID (Datatype: String)	Role (Datatype: String)
RADIUS			
alphanumeric	only plain text password supported	Where to get the data from? Tenant ID should be the ID of the Tenant created.	Where to get the data from? "System-Admin", "Tenant-Admin", "Tenant-Operator " are the only sup-ported roles, they must appear under the each user. Which vendor attribute is used to identify the role? User-Role attribute is used to iden-tify the role (see details below on how to configure the radius server)

Additional Information on LDAP

```
version: 1
dn: cn=Network-Operator,dc=ibm
objectclass: top
objectclass: groupOfNames
cn: Tenant-Operator
member: uid=varun1,dc=ibm

dn: uid=varun1,dc=ibm
objectclass: top
objectclass: inetOrgPerson
objectclass: person
objectclass: organizationalPerson
cn: varun
sn: tayur
ou: tenant1Id
uid: varun1
userPassword::
e1NTSEF9SmU4RGhETlhTK1JCuk42wwVhYURkU01zdZFjZF1Zdw1wUE1rYXc9P
Q==

dn: cn=Network-Admin,dc=ibm
objectclass: top
objectclass: groupOfNames
cn: Tenant-Admin
member: uid=varun2,dc=ibm
dn: dc=ibm
objectclass: top
objectclass: domain
dc: ibm
dn: uid=varun2,dc=ibm
objectclass: top
objectclass: inetOrgPerson
objectclass: person
objectclass: organizationalPerson
cn: varun
sn: tayur
ou: tenant2Id
uid: varun2
```

Configuring RADIUS Server

Radius server return Role and Tenant mapping to SDN VE components. The following vendor-specific attributes need to be specified on the Radius server.

1. Create a file under the free radius installation
<FR_INSTALL>\etc\raddb\dictonaries\dictionary.ibm
using the following file contents:

```
VENDOR Example 16122
#
# Standard attribute
#
BEGIN-VENDOR Example
ATTRIBUTE User-Role 1 string
ATTRIBUTE Tenant 2 string
END-VENDOR Example
```

2. Add an entry of the file we created in Step 1 in `<FR_INSTALL>\etc\raddb\dictionary:`

```
#
# This is the master dictionary file, which references
# the pre-defined dictionary files included with the
# server.
#
# Any new/changed attributes MUST be placed in this
# file, as
# the pre-defined dictionaries SHOULD NOT be edited.
#
# $Id$
#

#
# The filename given here should be an absolute path.
#
$INCLUDE dictionaries/dictionary
$INCLUDE dictionaries/dictionary.ibm
```

3. Edit the user information in `<FR_INSTALL>\etc\raddb\users:`

```
testuser ClearText-Password := "testpw"
        Reply-Message = "Hello, %{User-Name}",
        User-Role = "System-Admin",
        Tenant = "2"
```

Managing LDAP Server

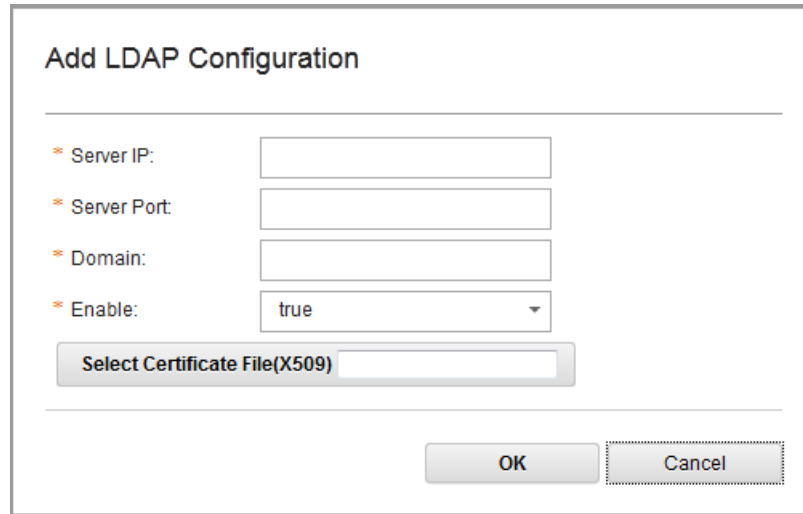
This section includes information on adding, modifying, and deleting LDAP server information on the IBM Unified Controller GUI.

Adding LDAP Server

On the GUI:

1. Select **Administration > System Tools > LDAP**.
2. Double-click on **LDAP**. The **LDAP Settings** page is displayed in the right pane.
3. Select **Add Config**.

4. Specify the Server IP, Server Port, Domain, and Enable status. Select the certificate file if NIST is enabled.



The image shows a dialog box titled "Add LDAP Configuration". It contains four labeled input fields: "Server IP:", "Server Port:", "Domain:", and "Enable:". The "Enable:" field is a dropdown menu currently set to "true". Below these fields is a button labeled "Select Certificate File(X509)" followed by a text input field. At the bottom right of the dialog are two buttons: "OK" and "Cancel".

5. Select **OK**.
If successfully added, the "LDAP configuration added successfully " is displayed on the top left corner.

Modifying Domain Name

The LDAP Server domain name can be modified as follows:

On the GUI:

1. Select **Administration > System Tools > LDAP**.
2. Double-click on **LDAP**. The **LDAP Settings** page is displayed in the right pane.
3. Select the LDAP server.
4. Select **Modify Domain**.
5. Enter the new domain name.
6. Select **OK**.

If successfully modified, the "LDAP configuration saved successfully "message is displayed on the top left corner.

Enabling/Disabling LDAP Service

On the GUI:

1. Select **Administration > System Tools > LDAP**.
2. Double-click on **LDAP**. The **LDAP Settings** page is displayed in the right pane.
3. Select the LDAP server.
4. Select **On/Off**.

The message "LDAP configuration saved successfully " is displayed on the top left corner.

The status in the "Enabled" column in "LDAP Info" table also toggles between true and false.

Deleting LDAP Configuration

On the GUI:

1. Select **Administration > System Tools > LDAP**.
2. Double-click on **LDAP**. The **LDAP Settings** page is displayed in the right pane.
3. Select the LDAP server.
4. Select **Delete Config**.

If successfully deleted, the “LDAP configuration deleted successfully “ message is displayed on the top left corner.

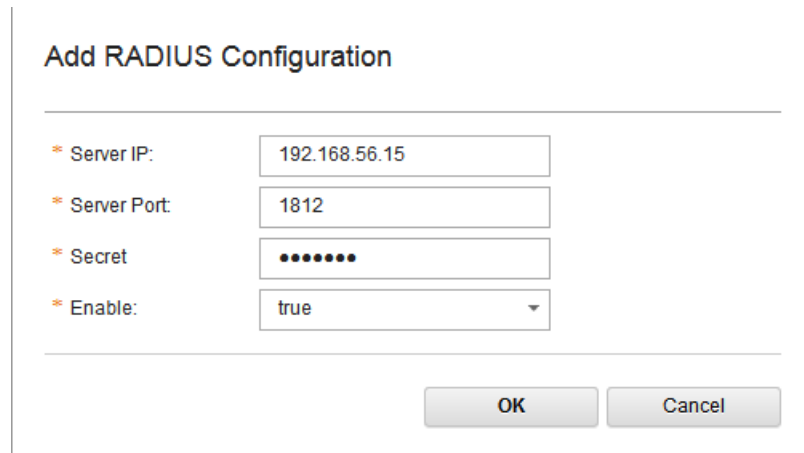
Managing RADIUS Server

This section includes information on adding, modifying, and deleting RADIUS server information on the IBM Unified Controller GUI.

Adding RADIUS Server

On the GUI:

1. Select **Administration > System Tools > RADIUS**.
2. Double-click on **RADIUS**. The **RADIUS Settings** page is displayed in the right pane.
3. Select **Add Config**.
4. Specify the Server IP, Server Port, Domain, and Enable status.



The image shows a dialog box titled "Add RADIUS Configuration". It contains four fields, each with a red asterisk icon indicating a required field:

- Server IP:** A text input field containing "192.168.56.15".
- Server Port:** A text input field containing "1812".
- Secret:** A password input field with a masked value represented by seven dots.
- Enable:** A dropdown menu with "true" selected.

At the bottom right of the dialog box are two buttons: "OK" and "Cancel".

5. Select **OK**.

If successfully added, the “RADIUS configuration added successfully “ is displayed on the top left corner.

Modifying Password

The RADIUS Server password can be modified as follows:

On the GUI:

1. Select **Administration > System Tools > RADIUS**.
2. Double-click on **RADIUS**. The **RADIUS Settings** page is displayed in the right pane.
3. Select the RADIUS server.
4. Select **Modify Secret**.
5. Enter the new password.
6. Select **OK**.

If successfully modified, the “RADIUS configuration saved successfully” message is displayed on the top left corner.

Enabling/Disabling RADIUS Service

On the GUI:

1. Select **Administration > System Tools > RADIUS**.
2. Double-click on **RADIUS**. The **RADIUS Settings** page is displayed in the right pane.
3. Select the RADIUS server.
4. Select **On/Off**.

The message “RADIUS configuration saved successfully” is displayed on the top left corner.

The status in the “Enabled” column in “RADIUS Info” table also toggles between true and false.

Deleting RADIUS Configuration

On the GUI:

1. Select **Administration > System Tools > RADIUS**.
2. Double-click on **RADIUS**. The **RADIUS Settings** page is displayed in the right pane.
3. Select the RADIUS server.
4. Select **Delete Config**.

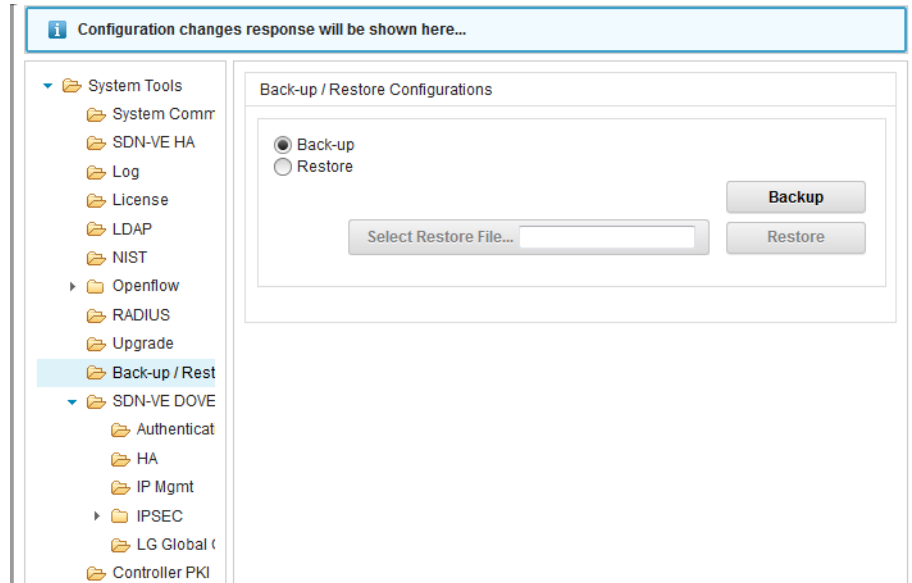
If successfully deleted, the “RADIUS configuration deleted successfully” message is displayed on the top left corner.

Managing Configuration

This section provides information on backing up and restoring SDN VE configuration using the IBM Unified Controller GUI.

On the GUI:

1. Select **Administration > System Tools > Back-up/Restore**.
2. Double-click on **Back-up/Restore**. The **Back-up/Restore Configurations** page is displayed in the right pane.



Backup Configuration

Before proceeding with backup, ensure the current configuration is saved.

1. Select **Back-up** radio button.
2. Select **Backup**.

The configuration is backed up locally as a `.tar` file.

Restore Configuration

Note: You must restore configuration only on a new VM i.e. the VM must not have any existing controller configuration. Perform a new installation of the controller and then restore. Performing a restore on a controller with existing user configuration will have unpredictable results. The restore operation will result in a restart of the controller.

The restore function can be used to clone a previously backed up configuration on a new controller installation.

For a controller cluster, follow step 1 to step 7.

For a standalone controller, follow step 2 to step 5.

Restore configuration as follows:

1. Configure HA. See [“Establish Unified Controller High-Availability” on page 41](#)
2. Power off the Standby controller.
3. Select **Administration > System Tools > Back-up/Restore > Restore** radio button on the Active controller.
4. Click on the **Select Restore File** dialog box to locate the file.
5. Select **Restore**. As part of restore, the Active controller will be restarted.
If successfully restored, the “Restore configurations completed successfully.” message is displayed on the top left corner.
6. Ensure you are able to login to the Active controller.
7. Power up the Standby controller.

Part 2: Advanced Features

Chapter 9. OpenStack

OpenStack is an open-source cloud computing platform deployed as an Infrastructure as a service (IaaS) solution. OpenStack contributes towards various components such as compute, storage, and networking.

The IBM SDN VE solution supports the OpenStack Neutron API, which is a network abstraction that allows OpenStack to use the underlying network as the infrastructure without requiring it to have knowledge of the underlying resources.

Note: If you wish to use OpenStack Neutron APIs for your SDN VE setup, you must only use OpenStack for setup, configuration, and implementation.

OpenStack Neutron API requires the use of a plugin to implement the logical API requests.

This chapter describes the installation of the IBM SDN VE plugin that integrates with the controller.

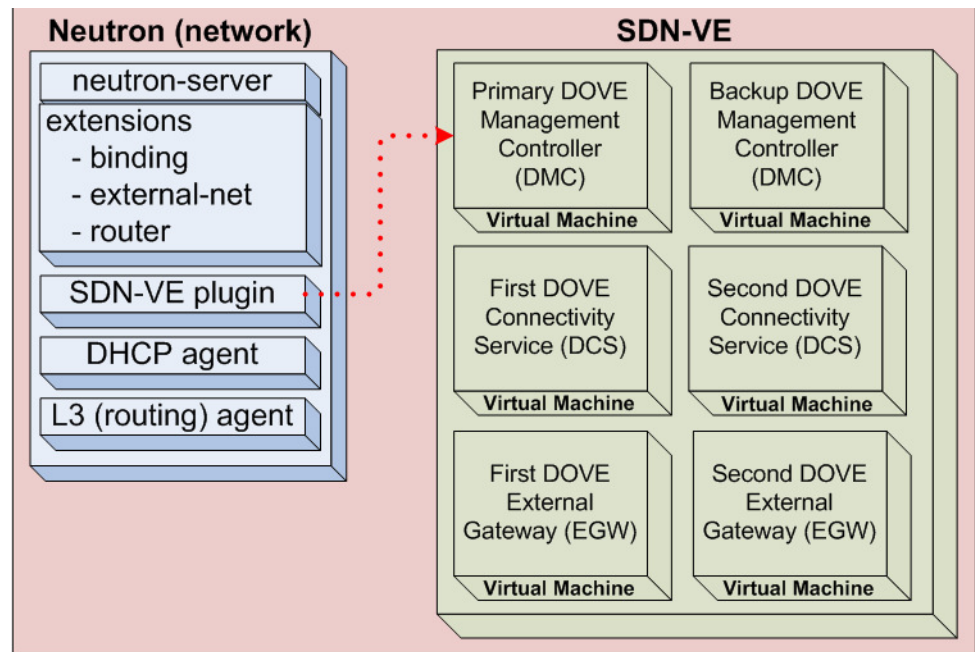


Figure 3. IBM SDN VE Plugin Integration

OpenStack Integration with SDN VE Plugin

Following instructions assume OpenStack with neutron is installed.

The SDN VE plugin configuration must be performed on the host that has the OpenStack Controller installed i.e. the controller that runs the neutron-server service.

Plugin Integration

Follow the steps in this section to integrate the SDN VE plugin with OpenStack.

You need OpenSSL version 1.0.1e-16 or higher installed on the host.

1. Remove Open vSwitch Neutron service agent:
 - a. Stop the agent (if running):

```
[root@sdnve-system2 ~]# service neutron-openvswitch-agent stop
Stopping neutron-openvswitch-agent: [ OK ]
[root@sdnve-system2 ~]#
```

- b. Remove the agent registration from Linux service framework:

```
[root@sdnve-system2 ~]# chkconfig --del neutron-openvswitch-agent
```

- c. Ensure the service file is available in the `/etc/rc` directory:

```
[root@sdnve-system2 ~]# find /etc/rc.d/ -name
"*neutron-openvswitch-agent*"
/etc/rc.d/init.d/neutron-openvswitch-agent
[root@sdnve-system2 ~]
```


2. Install or Upgrade the Plugin:

Untar the plugin file:

```
tar -zxvf sdnve-neutron-plugin-for-icehouse-only.tgz
```

The following directories are created with the relevant files placed in the directory:

```
sdnve-plugin-icehouse/  
sdnve-plugin-icehouse/plugin-archive/  
sdnve-plugin-icehouse/plugin-archive/int_support/  
sdnve-plugin-icehouse/plugin-archive/int_support/rc/  
sdnve-plugin-icehouse/plugin-archive/int_support/rc/neutron-sdnve-agent  
sdnve-plugin-icehouse/plugin-archive/int_support/sdnve_neutron_plugin  
.ini  
sdnve-plugin-icehouse/plugin-archive/int_support/neutron-sdnve-agent  
sdnve-plugin-icehouse/plugin-latest/  
sdnve-plugin-icehouse/plugin-latest/scripts/sdnve_plugin_install.py  
sdnve-plugin-icehouse/plugin-latest/.project  
sdnve-plugin-icehouse/plugin-latest/ibm/  
sdnve-plugin-icehouse/plugin-latest/ibm/sdnve_neutron_plugin.py  
sdnve-plugin-icehouse/plugin-latest/ibm/README  
sdnve-plugin-icehouse/plugin-latest/ibm/sdnve_api_fake.py  
sdnve-plugin-icehouse/plugin-latest/ibm/sdnve_api.py  
sdnve-plugin-icehouse/plugin-latest/ibm/__init__.py  
sdnve-plugin-icehouse/plugin-latest/ibm/common/  
sdnve-plugin-icehouse/plugin-latest/ibm/common/exceptions.py  
sdnve-plugin-icehouse/plugin-latest/ibm/common/constants.py  
sdnve-plugin-icehouse/plugin-latest/ibm/common/config.py  
sdnve-plugin-icehouse/plugin-latest/ibm/common/__init__.py  
sdnve-plugin-icehouse/plugin-latest/ibm/agent/  
sdnve-plugin-icehouse/plugin-latest/ibm/agent/sdnve_neutron_agent.py  
sdnve-plugin-icehouse/plugin-latest/ibm/agent/__init__.py
```

The Plugin file can be installed or upgraded using a script as follows:

```
[root@sdnve-system2 ~]# cd plugin-latest/scripts  
[root@sdnve-system2 scripts]#  
python sdnve_plugin_install.py {install | upgrade}  
[root@sdnve-system2 scripts]# cd /
```

To manually install or upgrade the plugin, proceed with steps a - f.

a. Copy the extracted file to the required directories using the commands:

```
[root@sdnve-system2 ~]# cp  
plugin-archive/int_support/neutron-sdnve-agent /usr/bin/  
[root@sdnve-system2 ~]# cp  
plugin-archive/int_support/rc/neutron-sdnve-agent /etc/init.d/
```

b. Create a directory as follows:

```
[root@sdnve-system2 ~]# mkdir -p /etc/neutron/plugins/ibm/
```

c. Copy the plugin file to the `ibm` directory:

```
[root@sdnve-system2 ~]# cp
plugin-archive/int_support/sdnve_neutron_plugin.ini
/etc/neutron/plugins/ibm/
```

d. Link the plugin file:

```
[root@sdnve-system2 ~]# ln -sf
/etc/neutron/plugins/ibm/sdnve_neutron_plugin.ini
/etc/neutron/plugin.ini
```

e. (If Upgrading) Move the files from
`/usr/lib/python2.6/site-packages/neutron/plugins/ibm`
directory to `/ibm.old` directory:

```
[root@sdnve-system2 ~]# mv
/usr/lib/python2.6/site-packages/neutron/plugins/ibm
/usr/lib/python2.6/site-packages/neutron/plugins/ibm.old
```

f. Copy the plugin files to the `/plugins` directory:

```
[root@sdnve-system2 ~]# cp -r plugin-latest/ibm
/usr/lib/python2.6/site-packages/neutron/plugins/
```

3. Register the plugin agent:

```
[root@sdnve-system2 ~]# chkconfig --add neutron-sdnve-agent
```

Verify the registration:

```
[root@sdnve-system2 ~]# find /etc/ -name "*neutron-sdnve-agent*"

/etc/rc.d/rc5.d/K02neutron-sdnve-agent
/etc/rc.d/rc1.d/K02neutron-sdnve-agent
/etc/rc.d/rc6.d/K02neutron-sdnve-agent
/etc/rc.d/rc3.d/K02neutron-sdnve-agent
/etc/rc.d/rc4.d/K02neutron-sdnve-agent
/etc/rc.d/init.d/neutron-sdnve-agent
/etc/rc.d/rc2.d/K02neutron-sdnve-agent
/etc/rc.d/rc0.d/K02neutron-sdnve-agent
[root@sdnve-system2 ~]#
```

4. Configure the plugin:

a. Edit the `sdnve_neutron_plugin.ini` file:

```
[root@sdnve-system2 ~]# vi
/etc/neutron/plugins/ibm/sdnve_neutron_plugin.ini

[sdnve]
integration_bridge = br-int
#use_fake_controller = True
#interface_mappings = default:eth2

#Provide comma separated controller IP/s
controller_ips=1.2.3.4
#userid=
#password=
default_tenant_type=OVERLAY

[agent]
root_helper = sudo /usr/bin/neutron-rootwrap
/etc/neutron/rootwrap.conf
# Agent's polling interval in seconds
# polling_interval = 2

[securitygroup]
# Firewall driver for realizing neutron security group function.
firewall_driver = neutron.agent.firewall.NoopFirewallDriver
# firewall_driver = neutron.agent.firewall.NoopFirewallDriver
# Example: firewall_driver =
neutron.agent.linux.iptables_firewall.OVSHybridIptablesFirewallDriver

[database]
#change ovs to sdnve below
sql_connection = mysql://neutron:neutron@127.0.0.1/neutron
```

In the `[sdnve]` section:

- Assign the SDN VE controller External IP address to `controller_ips` variable:

```
controller_ips=9.70.29.97
```

- (Optional) Change the admin password:

Delete the #

Assign the required password:

```
password=
```

In the `[database]` section:

- Specify the neutron database name:

```
sql_connection = mysql://neutron:neutron@127.0.0.1/neutron
```

Note: You can find out the database name using the following command:

```
[root@sdnve-system2 ~]# mysql -uneutron -pneutron -s -e  
'show databases' | grep -v -e "Database" -e  
"information_schema"  
ovs_neutron
```

The default login credentials for the database are:

username: neutron

password: neutron

5. Edit the neutron.conf file to specify the SDN VE plugin as the core plugin:

```
[root@sdnve-system2 ~]# vi /etc/neutron/neutron.conf  
  
core_plugin = neutron.plugins.ibm.sdnve_neutron_plugin.SdnvePluginV2  
  
#service_plugins =  
neutron.services.loadbalancer.plugin.LoadBalancerPlugin,neutron.serv  
ices.metering.metering_plugin.MeteringPlugin
```

6. Edit the Nova configuration file:

This edit is required on all OpenStack Nova compute hosts

```
[root@sdnve-system2 ~]# vi /etc/nova/nova.conf  
  
VIF driver : Generic VIF Driver  
Comment:: linuxnet_interface_driver and security_group_api  
  
libvirt_vif_driver = nova.virt.libvirt.vif.LibvirtGenericVIFDriver  
#linuxnet_interface_driver =  
#security_group_api = neutron
```

7. Restart the following OpenStack services:

```
[root@sdnve-system2 ~]# service neutron-server restart
Stopping neutron: [ OK ]
Starting neutron: [ OK ]

[root@sdnve-system2 ~]# service neutron-sdnve-agent restart
Stopping neutron-sdnve-agent: [FAILED]
Starting neutron-sdnve-agent: [ OK ]

[root@sdnve-system2 ~]# service openstack-nova-api restart
Stopping openstack-nova-api: [ OK ]
Starting openstack-nova-api: [ OK ]

[root@sdnve-system2 ~]# service openstack-nova-conductor restart
Stopping openstack-nova-conductor: [ OK ]
Starting openstack-nova-conductor: [ OK ]

[root@sdnve-system2 ~]# service openstack-nova-compute restart
Stopping openstack-nova-compute: [ OK ]
Starting openstack-nova-compute: [ OK ]

[root@sdnve-system2 ~]# service openstack-nova-scheduler restart
Stopping openstack-nova-scheduler: [ OK ]
Starting openstack-nova-scheduler: [ OK ]
```

Chapter 10. Waypoint Connectivity Service

A middlebox is a network appliance that resides between the source and destination of a packet. Typical middlebox examples include firewalls, Network Address Translators (NAT), load balancers, and Intrusion Prevention Systems (IPS)/Intrusion Detection Systems (IDS).

The IBM SDN VE solution supports routing of traffic through such middleboxes. Use of middleboxes for routing is also known as service insertion or Waypoint service enablement.

The SDN VE Unified Controller supports Waypoint service by enabling external devices performing such functions in your logical groups. The IBM SDN VE solution provides both transparent and non-transparent Waypoint services. Non-transparent services include routed NAT mode and routed mode.

Waypoint Service Operation

Transparent Mode

Waypoint devices operating in transparent mode do not change the MAC address information in the packet header. They primarily operate as a bridge for the packet to reach its destination. Examples of Waypoint devices that operate in transparent mode include firewalls, transparent proxy devices, and IPS/IDS.

Transparent Waypoint devices are not part of any particular subnet. The SDN VE Unified Controller uses the connectivity service configuration to route the packet to the appropriate Waypoint devices, and ultimately to the packet's destination. The controller redirects the packet to the ingress port of the Waypoint device. When the packet exits from the egress port of the Waypoint device, it is routed to the next destination.

Routed NAT Mode

Waypoint devices operating in routed NAT mode are capable of terminating incoming and establishing new connections with servers and clients. Load balancers, and the Direct Server Return method of load balancing, typically work on this model.

The SDN VE Unified Controller redirects packets to the Waypoint device's ingress interface. At this point, the source IP address/MAC address in the packet header is modified with the device's interface IP address, and the destination IP is the IP address of the intended recipient. Thus, the Routed NAT Waypoint terminates the connection from the original source and creates a new connection with the destination.

The SDN VE connectivity service is responsible for ensuring Layer 2 and Layer 3 connectivity between the Waypoint device and end stations. The connectivity service uses policies and Layer 3 routing to provide this service.

Additionally, Direct Server Return method of load-balancing can be supported.

Routed Mode

Waypoint devices operating in routed mode perform operations similar to a router: packets sent by an end-station VM will have their destination MAC address modified to the Waypoint device's ingress interface MAC address.

The SDN VE connectivity service is responsible for ensuring Layer 2 and Layer 3 connectivity between the Waypoint device and end stations. The connectivity service uses policies and Layer 3 routing to provide this service. The SDN VE connectivity service views these Waypoint devices as Explicit or Implicit devices. Based on the specified interface on which this service is applied, the controller uses appropriate logical routing to enable the connectivity.

Routed Explicit Devices

The Waypoint devices are configured to be the default gateway device for an end station Overlay VM. The controller sends all routed traffic to the MAC address of the Waypoint device. The device then logically forwards the packet out of its egress interface (configured as next hop in the connectivity service chain definition) to the destination network.

If a particular instance of the Waypoint device is modified, or if the device configuration changes, the end station needs to be reconfigured as required.

Some firewalls and Web proxy servers operate on this model.

Routed Implicit Devices

If a Waypoint device is configured as routed implicit, its existence is not known to the end stations. The implicit routing of packets to and from the Waypoint device provides the connectivity. Since an end station need not know about the Waypoint device, it does not need a reconfiguration if the connectivity instance of the Waypoint device is modified, or if the device configuration changes.

Some load balancers, firewalls, and Web proxy servers can be made to operate on this model.

Waypoint Connectivity

Waypoint devices use a single interface (one arm operation) to receive and send packets. Waypoint service can be configured by specifying the IP address or MAC address of the interface.

When you create a Waypoint device (middle box), a Connectivity Group (CG) is created. You must connect the Waypoint device.

If a Waypoint device—operating in transparent mode—has multiple interfaces, each interface must be placed in a separate CG. The SDN VE 5000V Distributed vSwitch uses these CGs to identify the traffic source, and then decides the next course of action.

For example, a Waypoint device W1 has two interfaces defined with two CGs: W11 and W12. Two service chains are defined as follows:

S1 = {W11, W12}

S2 = {W12, W11}

When S1 is applied as a policy between CGs E1 and E2, traffic from E1 will ingress the Waypoint on W11. The Waypoint device sends traffic to the SDN VE 5000V Distributed vSwitch out from W12 CG interface. The vSwitch sees that the traffic has come from W12 CG interface. Based on this, it determines the original source i.e. E1's CG, and then applies the correct service chain and forwards the packet to a VM in E2.

Waypoint devices operating in Routed mode and Routed NAT mode, can have multiple interfaces with the same CG. Traffic is forwarded based on the MAC addresses or IP addresses (if implicit gateways are used).

Waypoint Discovery

You must export the VNID of the connectivity groups and middle boxes you create to the virtual switch. The virtual switch uses this VNID to correlate the Waypoint device and connectivity group with the resource specified in the template.

You must create a Waypoint VM and assign a port profile with a VNID. You must create a Waypoint VNID and export it to the SDN VE vSwitch hosting the Waypoint. The Waypoint's NIC should be connected to this VNID bridge. The 5000VSDN VE vSwitch uses this VNID to correlate the Waypoint device with the resource specified in the template.

To export a VNID:

1. Login to the controller GUI and select the tenant.
2. Select **Services > Logical Groups**. The groups are displayed in the right pane menu.
3. Select the group name. The configured connectivity groups are displayed in the right pane display area.
4. Right-click on the connectivity group and select **Export VNID**. The **Export VNID** window is displayed.
5. Specify the vSwitch IP address.
6. Select **OK**.

Waypoint devices cannot be shared between tenants. To use the same Waypoint device across tenants, you must define new CGs per tenant and connect the Waypoint device's interface to these CGs.

Waypoint Configuration

Note: Waypoint configuration can be completed using the controller GUI. Command-line interface is not supported.

Note: You may also configure the Waypoint functionality using service templates and REST APIs for configuring middleboxes, service chains, and Waypoint policies.

Typically, the Waypoint device itself is not the destination for data traffic; it is part of a chain of Waypoint devices that carries traffic originating within one Connectivity Group and is destined for another Connectivity Group.

Waypoint device configuration includes:

- Defining a connectivity instance.
- Providing middlebox specifications.
- Configuring a service chain.

- Defining policies.

Waypoint Configuration Using Service Templates and REST APIs

Table 13 provides a description of the elements that need to be specified during the Waypoint configuration process. The “CRUD” designation in the table signifies which operation an element can be used in: Create, Read, Update, or Delete.

Table 13. Waypoint Configuration Entities

item	Description
Service Type	Type of virtual middlebox service and its general characteristics. Example: firewall, IPS, load balancer, Web gateway, ACL rule set
Service Instance	Middlebox instance configuration Specifications: firewall rules, load-balancing algorithm, traffic manipulation characteristics
Connectivity Instance	A connectivity pattern with specific virtual end points including networks and ports, as well as service instances.
Middlebox Parameters ^a	
Device name (CRUD)	Name; character string
Service type(CRD)	Firewall, loadbalancer, IPS; character string
connectivity type(CR)	Transparent, Routed, routedNAT
Data Interface(CRUD)	Ingress Port Group ID Egress Port Group ID
HA mode(CR)	Active-Active, Active-Standby, None
Mgmt Interface(CRUD)	IP address of management interface (if out-of-band mechanism is available and used)
Properties(CR)	required; yes/no
Tenant(CR)	Tenant ID

a. CRUD - Create, Read, Update; Delete

Defining a Connectivity Instance

A connectivity instance can be defined using a service template. The service template can be built using either the SDN VE Controller GUI, or by importing a predefined template file and tailoring it as per your requirement. You may also use the REST API service template defined in JSON.

Service templates include the following information:

- Resources: instances of OpenStack elements
 - unique name, resource type, optional/required properties
- Parameters: defined values that can be overridden at runtime
 - strings, numbers (with constraints), lists
 - dereference in Resources or Outputs sections
- Mappings
 - define lookup table as key-attr pairs
 - use `Fn::FindInMap` to get values
- Outputs
 - declare info to be passed back to user about an existing stack

Service Template Example

This section includes example service template. You can use this template and modify the information based on your requirement.

Example: Routed Load Balancer

Deployment template for routed load balancer:

```
{
  "HeatTemplateFormatVersion" : "0.1",
  "Description" : "GeneratedFrom Policy Chain
:287b4de3-4964-3eb0-9db5-fdalef80dbdb",
  "Resources" : {
    "2" : {
      "Type" : "OS::Quantum::Net",
      "Properties" : {
        "name" : "2"
      }
    },
    "9c1ae62a-1739-33d4-bcfd-152d381bc461" : {
      "Type" : "OS::Neutron::connectivity::service",
      "Properties" : {
        "ha_mode" : "NONE",
        "interface_type" : "one_arm",
        "service_type" : "loadbalancer",
        "name" : "f5-int",
        "required" : "yes",
        "health_check" : "false",
        "mode" : "routed_nat"
      }
    },
    "5" : {
      "Type" : "OS::Quantum::Net",
      "Properties" : {
        "name" : "5"
      }
    },
    "287b4de3-4964-3eb0-9db5-fdalef80dbdb" : {
      "Type" : "OS::Neutron::policy",
      "Properties" : {
        "policy_dest" : {
          "Ref" : "5"
        },
        "service_list" : [ {
          "Ref" : "9c1ae62a-1739-33d4-bcfd-152d381bc461"
        }, {
          "Ref" : "48f1fd3e-48a1-3525-a7f7-83e346e2cba2"
        } ],
        "policy_snaf_pool" : { },
        "policy_type" : "conn_service",
        "name" : "t1-p2-s1",
        "policy_src" : {
          "Ref" : "2"
        }
      }
    },
    "48f1fd3e-48a1-3525-a7f7-83e346e2cba2" : {
      "Type" : "OS::Neutron::connectivity::service",
      "Properties" : {
        "ha_mode" : "NONE",
        "interface_type" : "one_arm",
        "service_type" : "loadbalancer",
        "name" : "f5-ext",
        "required" : "yes",
        "health_check" : "false",
        "mode" : "routed_nat"
      }
    }
  }
}
```


After deploying a service template, you can view the connectivity instance on the controller GUI. Access the GUI using the following URL:

<https://<Controller HA external IPv4 address>:8443>

Waypoint Configuration Using Controller GUI

Providing Middlebox Specifications

The middlebox configuration can be specified on the controller GUI as follows:

1. Login to the controller GUI and select the tenant.
2. Select **Services > Connectivity Service**.
3. Select the **Middle Boxes** tab on the right pane.
4. Select the Create Middle Box icon ().
5. Specify the middle box service properties:

Service Properties

Name:

Connectivity Type:

TRANSPARENT

Service Type:

FIREWALL

HA Mode:

NONE

HA Service IP:

Required:

YES

Health Check

NO

OK

Cancel

Note: Fields that are not applicable to the current deployment are disabled.
Specify middle box details as follows:

Table 14. Middle Box Specifications

Component	Description
Name	Name of the middle box. The name must start with an alphabet but can have subsequent alphanumeric characters.
Connectivity Type	The type of middle box: Transparent Routed Routed_NAT
Service Type	The type of service provided by the middle box: Firewall Load Balancer DHCP IPS IDS NAT

Table 14. Middle Box Specifications

Component	Description
Health Check	<p>Enable or disable health check (for routed NAT device type only)</p> <p>If Health Check is enabled, a reverse policy is also deployed from the target connectivity group to the reverse NAT (RNAT) middlebox. If there are multiple such middleboxes in the chain, the reverse policies are defined from the target CG to the first such ROUTED_NAT device that has health check enabled.</p>
HA Node	<p>Select HA mode:</p> <p>Active-Active</p> <p>Active- Standby (only for Routed and Routed NAT devices)</p> <p>None</p>
HA Service IP	Required if the HA mode is Active-Standby
Required	<p>Controls the impact on the service chain if a middlebox goes offline. Values: YES or NO.</p> <p>If YES is selected: When middlebox goes offline, packets do not traverse to the next hop in the service chain.</p> <p>If NO is selected: When the middlebox goes offline, packets are sent to the next hop in the service chain.</p>

6. Select **OK**.

Configuring A Service Chain

A chain of Waypoint devices can be configured between the traffic source and destination. Whether traffic should pass through a Waypoint device or not can be specified using service chains and policies.

A set of Waypoint Connectivity Groups used to deploy service appliances are chained together in a sequence to form a service chain. You can apply the service chain as a policy between a pair of CGs that host endpoint VMs. A service chain can be applied to multiple CG pairs within a tenant. Traffic can be allowed or denied between two CGs; traffic can be diverted to flow through a series of Waypoint devices. A Waypoint device can be part of any service chain.

Note: Service chains can be used only for unicast traffic. For broadcast or multicast traffic, you can use a simple policy that allows or denies the traffic.

Note: Service chains cannot be configured between dedicated and shared groups.


Note: Service chains between shared groups can only be configured from the administrator tenant i.e. DOVE admin.

Service chains can be configured as follows:

On the controller GUI **Services > Connectivity Service** page:

1. Select the **Service Chains** tab on the right pane. You will see the list of configured middle boxes in the right column.
2. Select the required middle boxes. The selected middle boxes are displayed in the Service Chain display area.
3. Link the middle boxes as required.




4. Specify a service chain name and select the create icon ().

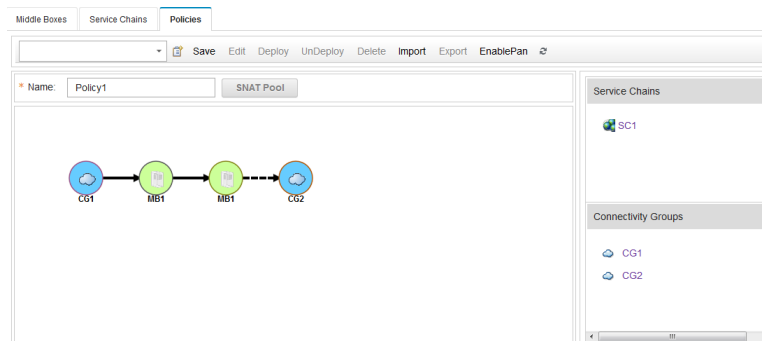
Defining a Policy

After configuring the service chain, you must define an end-to-end policy linking the service chain between two connectivity groups. at any point, only one service chain can be active in a policy.

Policies can be configured as follows:

On the controller GUI **Services > Connectivity Service** page:

1. Select the **Policies** tab on the right pane. You will see the list of configured service chains and connectivity groups in the right column.
2. Specify a policy name and select the create icon ().
3. Drag and drop the required service chains and connectivity groups in the Policies Chain display area. If the target connectivity group is of type external, SNAT address pool can be specified. See [“External Connectivity Groups - SNAT Pool Configuration” on page 144](#).
4. Link the service chains and connectivity groups as required.



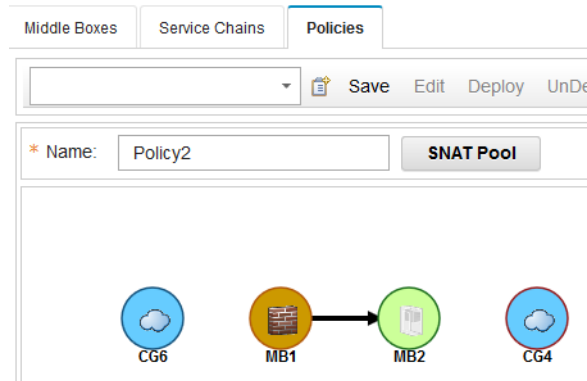
5. Select **Deploy** to activate the policy.

External Connectivity Groups - SNAT Pool Configuration

Connectivity groups can be configured with group type as dedicated, shared, or external.

When a connectivity group type is external, the connectivity group can communicate with external networks. You may assign a set of addresses for Network Address Translation. This is done while configuring the Connectivity Service Policies. See [“Defining a Policy” on page 143](#). If you do not specify the SNAT pool, the default pool configuration is used.

When the destination connectivity group (CG4 in this case) added in a policy configuration is of type external, the **SNAT Pool** button is activated.



Specify the NAT IP addresses and ports as follows:

1. Select **SNAT Pool**. The **SNAT Pool Configuration** window is displayed.

The 'SNAT Pool Configuration' window is displayed. It contains four input fields: 'Start IP', 'End IP', 'Start Port', and 'End Port'. At the bottom right, there are 'Save' and 'Cancel' buttons.

Specify the SNAT Pool details.

Table 15. SNAT Pool Specification

Component	Description
Start IP	Starting IP address of the range of addresses you want to allocate for NAT. The IP address range must be from the subnetwork to which the destination connectivity group belongs.
End IP	Ending IP address of the range of addresses you want to allocate for NAT.

Table 15. SNAT Pool Specification

Component	Description
Start Pool	Starting port number of the ports to be assigned for NAT. Port numbers can be in the range: 1-65535
End Pool	Starting port number of the ports to be assigned for NAT.

2. Select **Save**.

Waypoint High-Availability/Load Balancing

High-Availability and load balancing can be achieved by configuring multiple instances of a Waypoint device in the Waypoint Connectivity Group.

Transparent Mode

If the Waypoint device is operating in transparent mode, all VMs associated with the Waypoint CG are considered available for HA and load balancing purpose. However, each VM must be deployed on a separate SDN VE vSwitch since a vSwitch can host only one Waypoint instance at a point in time.

Routed/Routed NAT Mode

For Waypoint devices operating in Routed or Routed NAT mode, HA/load-balancing can be configured as Active/Active or Active/Stand-by.

In Active/Active type of HA/load-balancing, all VMs associated with a Waypoint CG are considered active. The connectivity service load-balances the flows to these devices by using a hash of {Source IP (SIP), Destination IP (DIP)}.

In Active/Stand-by type of HA/load-balancing, you must define a service IP address for the VM that is associated with the Waypoint CG. The VM for which the service IP address is configured will be considered to be active, and will receive the traffic flows. All other instances will be in Stand-by mode. The Active VM can forward traffic to other instances.

If a Waypoint instance goes down, traffic is redirected to another instance.

Limitations

- Waypoint devices configured in routed NAT mode replace the original IP of an incoming packet. These devices should not be shared between service chains because endpoints from one service chain may communicate with endpoints from another service chain because of the IP address replacement.
For example:
If a routed NAT Waypoint (Wnat) is used in two service chains (S3 and S4) as follows:
C1 → Wnat → S3 and C2 → Wnat → S4
Endpoints from C1 may be able to reach endpoints in S4, and endpoints in C2 may reach endpoints in S3.
To avoid this, configure two routed NAT Waypoints—one for each service chain—to ensure traffic is properly segregated.
- A middlebox cannot be edited if it is part of a service chain that belongs to a deployed policy. To edit such a middlebox, you must first remove the policy (**Controller GUI > Services > Connectivity Service > Policies > UnDeploy**).

- A middlebox cannot be deleted if it is part of a service chain.
- A middlebox can be added to multiple service chains. However, a middlebox can be added only once in the same service chain.
- Service Chain is an ordered list of middleboxes. This ordered list has to be unique. The same set of middleboxes in the same order cannot be added to multiple service chains. The service chains can have the same middleboxes in a different order, or a different set of middleboxes.
- A service chain cannot be edited if it is a part of a deployed policy. To edit such a service chain, you must first remove the policy (**Controller GUI > Services > Connectivity Service > Policies > UnDeploy**).
- A service chain cannot be deleted if it is part of a policy.
- Multiple policies that use Waypoint devices can be added between two connectivity groups. However, only one policy can be deployed at any point in time. This includes policies (without Waypoint devices) defined between two connectivity groups.
- External connectivity groups that are part of a policy as a source or target must have subnets associated with them. If it is not configured so, the policy will not deploy successfully. If the service chain has a routed NAT Waypoint, the connectivity groups connected with the routed NAT Waypoint must have a subnet associated with them.
- For external connectivity groups that are part of a policy, you must define a forwarding rule before deploying the policy. If a forwarding rule is not defined, the policy will not deploy successfully.
- A deployed policy cannot be deleted. To delete the policy, you must remove it first (**Controller GUI > Services > Connectivity Service > Policies > UnDeploy**).
- Only the IDs of connectivity groups that have been created can be specified in a HEAT template. The HEAT template may not import successfully if the IDs defined in the template and the IDs available on the controller do not match.

Chapter 11. NIST

The implementations specified in this section are compliant with National Institute of Standards and Technology (NIST) Special Publication (SP) 800-131A.

The IBM SDN VE controller can operate in NIST-compliant mode. By default, NIST is disabled on the controller.

If you enable NIST, the controller operates in strict mode: Encryption algorithms, protocols, and key lengths in strict mode are compliant with NIST SP 800-131A specification.

NIST mode ciphers—common to both Transport Layer Security (TLS) 1.1 and TLS 1.2 protocols—are used to ensure confidentiality of the data to and from the SDN VE components.

The following functions are compliant with NIST SP 800-131A specification:

- North Bound Communication: All north bound communication such as GUI, CLI, and REST APIs must be over HTTPS/TLSv1.1, TLSv1.2 using the acceptable cipher suites.
- South Bound Communication: All south bound communication over secure channel must happen using the acceptable cipher suites.
- HA Cluster: All nodes participating in a HA cluster must communicate over a secure channel using the acceptable cipher suites.
- User Authentication and Authorization: All external user authentication must be via secure LDAP using the acceptable cipher suites.
- Persistence: User passwords should be encrypted using AES 128 bit algorithm.
- External backup and restore: All external backup and restore must happen over HTTPS using the acceptable cipher suites.
- Product License Keys: Product license keys are encrypted using AES 128 bit encryption algorithm.
- Launching External GUI: Any external GUI launched within the SDN VE platform must be over HTTPS/TLSv1.1, TLSv1.2 using the acceptable cipher suites.

See [“Acceptable Cipher Suites” on page 148](#).

Enabling NIST

Note: NIST can be enabled on standalone IBM SDN VE controller or on controller nodes configured for high-availability (HA). In the latter case, the SDN VE Controller HA must be configured before enabling NIST. See [“Establish SDN VE Controller HA” on page 43](#).

Note: You must configure PKI controller private key, controller certificate, and CA root certificate before enabling NIST. See [“PKI Configuration” on page 151](#).

Enable NIST using the controller GUI as follows:

1. Login to the controller GUI: `https://<Controller HA external IPv4 address>:8443`.
2. Select **Administration > System Tools > NIST**.
3. Double-click on **NIST**. The **NIST Settings** page with the current status is displayed.
4. Select **Edit**.

5. Select **Enable** from the drop-down menu.
6. Select **Save**.

Note: You must run the system command `restart` on both the primary and secondary controllers every time you change the NIST setting. Restart the secondary controller only after the primary controller GUI comes up after the restart.

Acceptable Cipher Suites

The following cipher suites are acceptable (listed in the order of preference) as per the NIST SP 800-131A specification:

Table 16. List of Acceptable Cipher Suites in Strict Mode

Key Exchange	Encryption	MAC	Cipher Name
RSA	3DES_EDE_CBC	SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA
DH_DSS	3DES_EDE_CBC	SHA	TLS_DH_DSS_WITH_3DES_EDE_CBC_SHA
DH_RSA	3DES_EDE_CBC	SHA	TLS_DH_RSA_WITH_3DES_EDE_CBC_SHA
DHE_DSS	3DES_EDE_CBC	SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA
DHE_RSA	3DES_EDE_CBC	SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA
RSA	AES_128_CBC	SHA	TLS_RSA_WITH_AES_128_CBC_SHA
RSA	AES_256_CBC	SHA	TLS_RSA_WITH_AES_256_CBC_SHA
RSA	AES_128_CBC	SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256
RSA	AES_256_CBC	SHA256	TLS_RSA_WITH_AES_256_CBC_SHA256
DH_DSS	AES_128_CBC	SHA	TLS_DH_DSS_WITH_AES_128_CBC_SHA
DH_RSA	AES_128_CBC	SHA	TLS_DH_RSA_WITH_AES_128_CBC_SHA
DHE_DSS	AES_128_CBC	SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA
DHE_RSA	AES_128_CBC	SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA
DH_DSS	AES_256_CBC	SHA	TLS_DH_DSS_WITH_AES_256_CBC_SHA
DH_RSA	AES_256_CBC	SHA	TLS_DH_RSA_WITH_AES_256_CBC_SHA
DHE_DSS	AES_256_CBC	SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA
DHE_RSA	AES_256_CBC	SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA
DH_DSS	AES_128_CBC	SHA256	TLS_DH_DSS_WITH_AES_128_CBC_SHA256
DH_RSA	AES_128_CBC	SHA256	TLS_DH_RSA_WITH_AES_128_CBC_SHA256
DHE_DSS	AES_128_CBC	SHA256	TLS_DHE_DSS_WITH_AES_128_CBC_SHA256
DHE_RSA	AES_128_CBC	SHA256	TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
DH_DSS	AES_256_CBC	SHA256	TLS_DH_DSS_WITH_AES_256_CBC_SHA256
DH_RSA	AES_256_CBC	SHA256	TLS_DH_RSA_WITH_AES_256_CBC_SHA256
DHE_DSS	AES_256_CBC	SHA256	TLS_DHE_DSS_WITH_AES_256_CBC_SHA256
DHE_RSA	AES_256_CBC	SHA256	TLS_DHE_RSA_WITH_AES_256_CBC_SHA256

If using any of the following, please note the supported versions and configurations:

- Java: Java JDK 7.0 SR1

- Browsers: TLS 1.1 and TLS 1.2 are supported on the following browsers:
 - Firefox Version 24 or higher
 - On the URL field, type 'about:config'.
 - In the search field, type:
 - Security.tls.version.max; set the value to 3
 - Security.tls.version.min; set the value to 1
 - IE Version 10
 - Go to **Setting > Internet Options > Advanced > Security**
 - Enable "Use TLS 1.2", "Use TLS 1.1"
- CLI: Python with support for TLS 1.1 and TLS 1.2. The python library must be upgraded to support NIST-compliant cipher suites.

LDAP Configuration

Note: This section is required only if the NIST strict mode is enabled on the controller. See ["Enabling NIST" on page 147](#).

Configure LDAP as follows:

1. Download the trusted LDAP server certificate in PEM format.
2. Upload the certificate using the controller GUI:
 - a. Login to the controller GUI: `https://<Controller HA external IPv4 address>:8443`.
 - b. Select **Administration > System Tools > LDAP**.
 - c. Double-click on **LDAP**. The **LDAP Settings** page is displayed.
 - d. Select **Add Config**. The **Add LDAP Configuration** window is displayed.
 - e. Specify the LDAP server details.
 - f. Click the **Select Certificate File(X509)** field.
 - g. Browse and select an LDAP certificate, for LDAP client to trust.

The screenshot shows a window titled "Add LDAP Configuration". It contains four labeled input fields: "Server IP" with the value "9.121.62.106", "Server Port" with "3355", "Domain" with "cn=ibm", and "Enable" with a dropdown menu showing "true". Below these fields is a button labeled "Select Certificate File(X509)" followed by the text "server_cert.pem". At the bottom right of the window are two buttons: "OK" and "Cancel".

Note: Multiple certificates (in PEM format i.e. X509 format) can be uploaded in a single file separated by BEGIN CERTIFICATE and END CERTIFICATE BLOCKS.

- h. Select **OK**.

Chapter 12. Public Key Infrastructure

A Public Key Infrastructure (PKI) assures secure exchange of data using a public and a private cryptographic key pair. This key pair is exchanged via a trusted authority.

PKI includes the following:

- Certificate authority (CA): Issues and verifies digital certificates.
- Registration authority (RA): Verifies identity of the users/applications that request information from the CA.

The IBM SDN VE Controller and the Distributed Service Appliance (DSA) can be configured to use PKI. By default, security is enabled and authentication is disabled.

PKI Configuration

This section provides information on importing/uploading CA certificate, CRL certificate, and private key and certificate on the controller. You must generate the certificates and keys as per your PKI scheme.

Note: You must run the system command `restart` on both the primary and secondary controllers after you upload a certificate or key. Restart the secondary controller only after the primary controller GUI comes up after the restart.

1. Login to the controller GUI: `https://<Controller IP address>:8443`.
2. Select **Administration > System Tools > Controller PKI**.
The **Controller PKI Configuration** page is displayed.

3. Upload the appropriate certificates:

Key	Subject	Validity
EMAILADDRESS=admin@newcacerts.com, CN=Certification, OU=Certs CRL Division, O=New CA Certification Company Inc., L=Bangalore, ST=Karnataka, C=IN	EMAILADDRESS=admin@newcacerts.com, CN=Certification, OU=Certs CRL Division, O=New CA Certification Company Inc., L=Bangalore, ST=Karnataka, C=IN	true

Where:

- Controller private key: A controller private key in PEM (X509) format, used by the controller.
- Controller certificate: A controller certificate in PEM (X509) format, exchanged during TLS handshake.
- CA Root Certificate: Root Certificate(s) in PEM (X509) format, that are trusted by the controller.

Note: Multiple root certificates can be uploaded in a single file. The certificates can be separated by blocks of

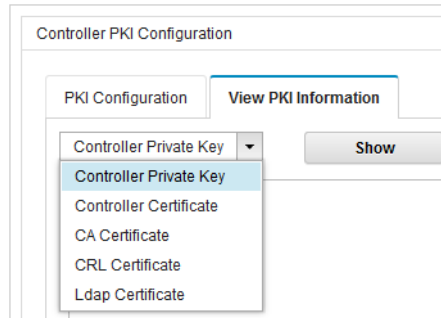
```
-----BEGIN CERTIFICATE-----
-----END CERTIFICATE-----
```

- CRL Verification list: A CRL verification list in PEM (X509) format.

Note: Multiple CRL verification lists can be uploaded in a single file. The lists can be separated by blocks of

```
-----BEGIN X509 CRL-----
-----END X509 CRL-----
```


4. View PKI information:
 - a. Select **View PKI Information** tab.
 - b. Select the Key/Certificate to view.



- c. Select **Show**.

DSA Configuration

PKI can be configured on the DSA using the DSA CLI. See [Chapter 22, “PKI Configuration Commands”](#) for the commands to upload certificates and keys.

5000V Agent Host Configuration

Configure the PKI settings on the 5000V Agent host as follows:

1. Ensure that the host VIB module is correctly installed:

```
~ # esxcli software vib list | grep ibm
ibm-esx-5000V                1.2.0-140704                IBM
VMwareAccepted              2014-08-06
~ #
```

2. Backup the existing (default) RUI files:

```
~ # cd /opt/ibm/sbin
sbin # cp rui.crt rui.crt.backup
sbin # cp rui.key rui.key.backup
```

3. Delete the original files:

```
sbin # rm rui.crt
sbin # rm rui.key
```

4. Copy your certificates and key to the current folder. Rename the files to the default names. See example:

```
sbin # mv server_cert.pem rui.crt
sbin # mv server_key.pem rui.key
```

5000V Controller Configuration

Use the following commands to import the PKI certificate and key files on the 5000V controller:

```
5000V(config)# ssl-import key <path and file name>
5000V(config)# ssl-import certificate <path and file name>
5000V(config)# ssl-import ca-root-cert <path and file name>
```


Use the following commands to enable CRL verification:

```
5000V(config)# security crl-enable
```

Deleting Certificates

You can delete certificates that you no longer need or are expired.

1. Select the **PKI Configuration** tab. You will see the list of certificates in the **CA Root Certificates** section.
2. Right-click on the certificate you want to delete.

CA Root Certificates 

Key	Subject	Validity
OU=STG, O="IBM India Pvt L Delete" L=Bangalore, ST=Karnataka, C=IN	OU=STG, O="IBM India Pvt Ltd.", L=Bangalore, ST=Karnataka, C=IN	true

3. Select **Delete**.

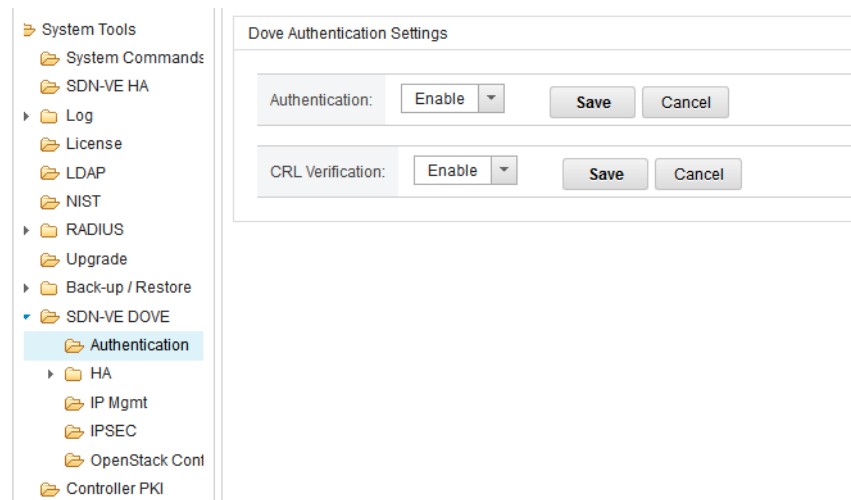
Authentication

Authentication is disabled by default. Authentication and certificate revocation list (CRL) can be configured for the SDN VE DOVE components as follows:

Note: You must run the system command `restart` on both the primary and secondary controllers every time you change the authentication setting. Restart the secondary controller only after the primary controller GUI comes up after the restart.

On the controller GUI:

1. Select **Administration > System Tools > SDN-VE DOVE > Authentication**. The **DOVE Authentication Settings** page is displayed.
2. Select **Edit to Enable/Disable Authentication and CRL Verification**.



Where:

- Enable Authentication: Authentication can be enabled between:
 - Controller and REST APIs/CLI client
 - Controller and SDN VE DOVE components
- Enable CRL verification between:
 - Controller and REST APIs/ CLI client
 - Controller and SDN VE DOVE components

3. Select **Save**.
4. Execute the following command on both the primary and secondary controllers.

Note: Restart the secondary controller only after the primary controller GUI comes up after the restart.

- a. Log in to the controller: `https://<Controller HA external IP address>:8443`.
 - b. Select **Administration > System Tools > System Commands**. The **System Commands** page is displayed.
 - c. Select **restart** from the **System Commands** drop-down list.
 - d. Select **Execute**.
5. Upload a CA certificate using any browser. Mozilla Firefox is used as an example here.
 - a. Open Mozilla Firefox.
 - b. Select **Tools > Options > Advanced > Certificates**.
 - c. Select **View Certificates > Your Certificates**.
 - d. Import the CA certificate. (See [“PKI Configuration” on page 151.](#))
 - e. Select **OK**.

You can now access the controller GUI.

Enabling Authentication on the DSA

Use the following DSA CLI command to enable authentication:

```
SDN-VE-DSA (config)# security-mode auth enable
```

Enable Authentication on the 5000V Controller

Use the following 5000V controller command to enable authentication:

```
5000V(config)# security client-auth-mode
```

IP Security

IP Security (IPSec) can be enabled between the SDN VE DOVE components configured for High-Availability (HA) i.e. the Primary and Secondary controllers. For IPSec service to function, authentication must be enabled (See [“Authentication” on page 154](#)) and PKI must be configured ([“PKI Configuration” on page 151](#)).

Note: HA must already be configured between the two controllers.

On the controller GUI:

1. Select **Administration > System Tools > SDN-VE DOVE > IPSEC**. The **IPSEC Settings** page is displayed.

The screenshot shows the IPSEC Settings page. On the left, a navigation tree lists various system tools, with 'IPSEC' highlighted under the 'SDN-VE DOVE' section. The main content area is titled 'IPSEC Settings'. It features an 'IPSEC Status' section with a refresh icon. Below this is a table with the following data:

Peer Ip:	NULL
Icmp State:	Not Established
Tcp State:	Not Established

At the bottom of the settings panel, there are two buttons: 'Enable IPSEC' and 'Disable IPSEC'.

2. Select **Enable IPSEC**.

Chapter 13. Access Control Lists

Access Control Lists (ACLs) are filters that permit or deny traffic for security purposes. Each filter defines the conditions that must match for inclusion in the filter, and also the actions that are performed when a match is made.

IBM Software Defined Network for Virtual Environments 1.2 supports the following ACLs:

- IPv4 Standard and Extended ACLs

Up to 127 ACLs are supported for networks that use IPv4 addressing. IPv4 ACLs are configured using the following command path:

```
iSwitch(config)# access-list ip <128-254> {standard|extended} ?
```

- MAC Extended ACLs

Up to 127 ACLs are supported for networks that use IPv4 addressing. MAC Extended ACLs are configured using the following command path:

```
iSwitch(config)# access-list mac extended <1-127> ?
```

ACLs can be applied only to traffic that ingresses a port or a profile group.

MAC Extended ACLs

MAC Extended ACLs use source and destination MAC addresses, along with optional protocol information, as the matching criteria. Up to 127 MAC Extended ACLs can be configured. These ACLs are numbered 1-127. MAC Extended ACLs have a higher priority than IPv4 ACLs.

IPv4 ACLs

IPv4 Standard ACLs use source and destination IPv4 addresses as the matching criteria.

IPv4 Extended ACLs use source and destination IPv4 addresses, along with optional protocol information, as the matching criteria.

Up to 127 IPv4 ACLs (Standard and Extended) can be configured. These ACLs are numbered 128-254.

Summary of Packet Classifiers

ACLs allow you to classify packets according to a variety of content in the packet header (such as the source address, destination address, source port number, destination port number, and others). Once classified, packet flows can be identified for more processing.

IPv4 ACLs and MAC Extended ACLs allow you to classify packets based on the following packet attributes:

- Ethernet header options
 - Source MAC address
 - Destination MAC address
 - VLAN number and mask
 - Ethernet type (ARP, IP, IPv6, MPLS, RARP, etc.)
 - Ethernet Priority (the IEEE 802.1p Priority)
- IPv4 header options (for IPv4 Standard ACLs)
 - Source IPv4 address and subnet mask
 - Destination IPv4 address and subnet mask

- IPv4 header options (for IPv4 Extended ACLs)
 - Source IPv4 address and subnet mask
 - Destination IPv4 address and subnet mask
 - IP protocol number or name as shown in [Table 17](#):

Table 17. Well-Known Protocol Types

Number	Protocol Name
1	icmp
2	igmp
6	tcp
17	udp
89	ospf
112	vrrp

- TCP/UDP header options
 - TCP/UDP application source port and mask as shown in [Table 18](#)
 - TCP/UDP application destination port as shown in [Table 18](#)

Table 18. Well-Known Application Ports

Port	TCP/UDP Application	Port	TCP/UDP Application	Port	TCP/UDP Application
20	ftp-data	79	finger	179	bgp
21	ftp	80	http	194	irc
22	ssh	109	pop2	220	imap3
23	telnet	110	pop3	389	ldap
25	smtp	111	sunrpc	443	https
37	time	119	nntp	520	rip
42	name	123	ntp	554	rtsp
43	whois	143	imap	1645/1812	Radius
53	domain	144	news	1813	Radius
69	tftp	161	snmp	1985	Accounting
70	gopher	162	snmptrap		hsrp

- TCP/UDP flag value as shown in [Table 19](#)

Table 19. Well-Known TCP flag values

Flag	Value
URG	0x0020
ACK	0x0010
PSH	0x0008
RST	0x0004
SYN	0x0002
FIN	0x0001

- ICMP message code and type as shown in [Table 20](#)

Table 20. Well-Known ICMP Messages and Codes

Description	ICMP Type	Code
Echo Reply	0	0
Destination Unreachable	3	0 = net unreachable 1 = host unreachable 2 = protocol unreachable 3 = port unreachable 4 = fragmentation needed and DF set 5 = source route failed
Source Quench	4	0
Redirect	5	0 = Redirect datagrams for the Network. 1 = Redirect datagrams for the Host. 2 = Redirect datagrams for the Type of Service and Network. 3 = Redirect datagrams for the Type of Service and Host.
Echo	8	0
Time Exceeded	11	0 = time to live exceeded in transit; 1 = fragment reassembly time exceeded.
Parameter Problem	12	0 = pointer indicates the error.
Timestamp	13	0
Timestamp Reply	14	0
Information Request	15	0
Information Reply	16	0

Summary of ACL Actions

Multiple ACLs can be applied to a port. Priority of the ACL is based on its ID; Lower IDs have higher priority and vice versa. When traffic ingresses a port, it is matched against the highest priority ACL. If no matching criteria are found, the next ACL is considered. This process continues until a match is found. If no match is found, traffic is permitted.

Assigning Individual ACLs to a Port

Once you configure an ACL, you must assign the ACL to the appropriate ports. Each port can accept multiple ACLs, and each ACL can be applied for multiple ports.

To assign an individual ACLs to a port, use the following IP Interface Mode commands:

IPv4 ACL

```
iSwitch(config)# interface port <port>  
iSwitch(config-if)# ip access-group <IPv4 ACL number> in
```

MAC Extended ACL

```
iSwitch(config)# interface port <port>  
iSwitch(config-if)# mac access-group <MAC ACL number> in
```

When multiple ACLs are assigned to a port, higher-priority ACLs are considered first, and their action takes precedence over lower-priority ACLs.

Assigning Individual ACLs to a VNIC Profile

You can assign ACLs to a VNIC profile. You can assign up to 254 ACLs (IPv4 and MAC ACLs together) to a VNIC profile. Each ACL can be applied for multiple VNIC profiles.

To assign an individual ACLs to a VNIC profile, use the following commands:

```
iSwitch(config)# iswitch vnicprof prof1 access-list mac <MAC ACL number>  
in  
iSwitch(config)# iswitch vnicprof prof1 access-list ip <IPv4 ACL number> in
```

When multiple ACLs are assigned to a VNIC profile, higher-priority ACLs are considered first, and their action takes precedence over lower-priority ACLs.

Viewing ACL Statistics

ACL statistics display how many packets have “hit” (matched) each ACL. Use ACL statistics to check filter performance or to debug the ACL filter configuration.

You must enable statistics for each ACL that you wish to monitor:

```
iSwitch(config)# access-list ip standard <ACL number> statistics  
iSwitch(config)# access-list ip extended <ACL number> statistics  
iSwitch(config)# access-list mac extended <ACL number> statistics
```

Statistics can be viewed using the following command:

```
iSwitch(config)# show access-list [<ACL number>] counters
```

Deleting ACLs

Use the following commands to delete an ACL:

IPv4 ACLs:

```
iSwitch(config)# no access-list ip <ACL number> {standard|extended}
```

MAC Extended ACLs:

```
iSwitch(config)# no access-list mac extended <ACL number>
```

ACLs Assigned to Profiles

ACLs assigned to ports that are part of a VNIC profile or DOVE profile cannot be deleted from ports using the commands mentioned in this section. The ACL must be deleted from the profile. Use the following commands:

VNIC Profile:

```
iSwitch(config)# iswitch vnicprof <profile name>  
iSwitch(config-vnicprof)# no access-list {mac <1-127>|ip <128-254>} in
```

DOVE Profile:

```
iSwitch(config)# iswitch doveprof <profile name>  
iSwitch(config-dvprof)# no access-list {mac <1-127>|ip <128-254>} in
```

ACL Configuration Examples

ACL Example 1

Use this configuration to block traffic to a specific host. All traffic that ingresses on port 1 is denied if it is destined for the host at IP address 100.10.1.1

1. Configure an Access Control List.

```
iSwitch(config)# access-list ip 150 standard  
iSwitch(config-std-nacl)# deny any host 100.10.1.1
```

2. Assign ACL 1 to port 1.

```
iSwitch(config)# interface port 1  
iSwitch(config-if)# ip access-group 150 in  
iSwitch(config-if)# exit
```

3. Verify configuration:

```
iSwitch(config)# show access-lists 150

Standard IP Access List 150
-----
Source IP address          : 0.0.0.0
Source IP address mask     : 0.0.0.0
Destination IP address     : 100.10.1.1
Destination IP address mask : 255.255.255.255
In Port List               : 1
Filter Action               : Deny
Status                     : InActive
```

ACL Example 2

Use this configuration to block traffic from a network destined for a specific host address. All traffic that ingresses in port 10 with source IP from class 100.10.1.0/24 and destination IP 200.20.2.2 is denied.

1. Configure an Access Control List.

```
iSwitch(config)# access-list ip 160 standard
iSwitch(config-std-nacl)# deny 100.10.1.0 255.255.255.0 host
200.20.2.2
iSwitch(config-std-nacl)# exit
```

2. Assign ACL 160 to port 10.

```
iSwitch(config)# interface port 10
iSwitch(config-if)# ip access-group 160 in
iSwitch(config-if)# exit
```

ACL Example 3

Use this configuration to block HTTP traffic on a port. All HTTP traffic that ingresses in port 12 is denied.

1. Configure an Access Control List.

```
iSwitch(config)# access-list ip 170 extended
iSwitch(config-ext-nacl)# deny tcp any any eq 80
iSwitch(config-ext-nacl)# exit
```

2. Assign ACL 170 to port 12.

```
iSwitch(config)# interface port 12
iSwitch(config-if)# ip access-group 170 in
iSwitch(config-if)# exit
```

ACL Example 4

Use this configuration to block all traffic except traffic of certain types. HTTP/HTTPS, DHCP, and ARP packets are permitted on the port. All other traffic is denied.

1. Configure an ACL for each type of traffic you want to permit.

```
iSwitch(config)# access-list ip 200 extended
iSwitch(config-ext-nacl)# permit tcp any any eq 80
iSwitch(config-ext-nacl)# exit
iSwitch(config)# access-list ip 210 extended
iSwitch(config-ext-nacl)# permit tcp any any eq 443
iSwitch(config-ext-nacl)# exit
iSwitch(config)# access-list ip 220 extended
iSwitch(config-ext-nacl)# permit udp any any eq 67
iSwitch(config-ext-nacl)# exit
iSwitch(config)# access-list ip 230 extended
iSwitch(config-ext-nacl)# permit udp any any eq 68
iSwitch(config-ext-nacl)# exit
```

2. Configure an ACL to deny all other traffic.

```
iSwitch(config)# access-list ip 240 extended
iSwitch(config-ext-nacl)# deny tcp any any
iSwitch(config-ext-nacl)# exit
iSwitch(config)# access-list ip 245 extended
iSwitch(config-ext-nacl)# deny udp any any
iSwitch(config-ext-nacl)# exit
```

Note: ACLs that permit traffic must have a higher priority than the ACLs that deny all traffic.

3. Configure a MAC ACL for each type of traffic that you want to permit. This example permits ARP traffic.

```
iSwitch(config)# access-list mac extended 10
iSwitch(config-ext-macl)# permit any any 806
iSwitch(config-ext-macl)# exit
```

4. Assign the ACLs to a port:

```
iSwitch(config)# interface port 7
iSwitch(config-if)# ip access-group 200 in
iSwitch(config-if)# ip access-group 210 in
iSwitch(config-if)# ip access-group 220 in
iSwitch(config-if)# ip access-group 230 in
iSwitch(config-if)# ip access-group 240 in
iSwitch(config-if)# ip access-group 245 in
iSwitch(config-if)# mac access-group 10 in
```

ACL Example 5

Use the following configuration to assign an ACL to a VNIC profile.

Note: When an ACL is added to a VNIC profile, the ACL is applied to all the ports that are part of the profile.

1. Create a VNIC profile: prof1.

```
iSwitch(config)# iswitch vnicprof create prof1
```

2. Configure ACLs for the VNIC profile.

```
iSwitch(config)# access-list mac extended 10
iSwitch(config-ext-macl)# permit any any ipv4
iSwitch(config-ext-macl)# exit

iSwitch(config)# access-list ip 245 extended
iSwitch(config-ext-nacl)# deny udp any any
iSwitch(config-ext-nacl)# exit
```

3. Add ACLs to the VNIC profile.

```
iSwitch(config)# iswitch vnicprof prof1 access-list mac 10 in
iSwitch(config)# iswitch vnicprof prof1 access-list ip 245 in
```

ACL Example 6

Use the following configuration to assign an ACL to a DOVE profile.

Note: When an ACL is added to a DOVE profile, the ACL is applied to all the ports that are part of the profile.

1. Export a network from the DOVE Management Console (DMC) to create a DOVE profile. For example: domain1.network1.vds1.
2. Configure ACLs for the DOVE profile.

```
iSwitch(config)# access-list mac extended 10
iSwitch(config-ext-macl)# permit any any ipv4
iSwitch(config-ext-macl)# exit

iSwitch(config)# access-list ip 245 extended
iSwitch(config-ext-nacl)# deny udp any any
iSwitch(config-ext-nacl)# exit
```

3. Add ACLs to the DOVE profile.

```
iSwitch(config)# iswitch doveprof domain1.network1.vds1
iSwitch(config-dvprof)# access-list mac 10 in
iSwitch(config-dvprof)# access-list ip 245 out
```

Chapter 14. Quality of Service

Quality of Service (QoS) features allow you to allocate network resources to mission-critical applications at the expense of applications that are less sensitive to such factors as time delays or network congestion. You can configure your network to prioritize specific types of traffic, ensuring that each type receives the appropriate QoS level.

The following topics are discussed in this section:

- [“QoS Overview” on page 167](#)
- [“Using DSCP Values to Provide QoS” on page 167](#)
- [“Using 802.1p Priority to Provide QoS” on page 168](#)

QoS Overview

QoS helps you allocate guaranteed bandwidth to the critical applications, and limit bandwidth for less critical applications. Applications such as video and voice must have a certain amount of bandwidth to work correctly; using QoS, you can provide that bandwidth when necessary. Also, you can put a high priority on applications that are sensitive to timing out or that cannot tolerate delay, by assigning their traffic to a high-priority queue.

By assigning QoS levels to traffic flows on your network, you can ensure that network resources are allocated where they are needed most. QoS features allow you to prioritize network traffic, thereby providing better service for selected applications.

The basic QoS model works as follows:

- Classify traffic:
 - Read DSCP value.
 - Read 802.1p priority value.

Using DSCP Values to Provide QoS

The IBM SDN VE 5000V Distributed vSwitch uses the Differentiated Services (DiffServ) architecture to provide QoS functions. DiffServ is described in IETF RFCs 2474 and 2475.


The six most significant bits in the TOS byte of the IP header are defined as DiffServ Code Points (DSCP). Packets are marked with a certain value depending on the type of treatment the packet must receive in the network device. DSCP is a measure of the Quality of Service (QoS) level of the packet.

The switch can classify traffic by reading the DiffServ Code Point (DSCP) or IEEE 802.1p priority value. When network traffic attributes match those specified in a traffic pattern, the policy instructs the controller to perform specified actions on each packet that passes through it.

Differentiated Services Concepts

To differentiate between traffic flows, packets can be classified by their DSCP value. The Differentiated Services (DS) field in the IP header is an octet, and the first six bits, called the DS Code Point (DSCP), can provide QoS functions. Each packet carries its own QoS state in the DSCP. There are 64 possible DSCP values (0-63).

Figure 4. Layer 3 IPv4 packet



The vSwitch can use the DSCP value to direct traffic prioritization.

QoS Levels

[Table 21](#) shows the default service levels provided by the vSwitch, listed from highest to lowest importance:

Table 21. Default QoS Service Levels


Service Level	802.1p Priority
Network Control	7
Internetwork Control	6
Voice	5
Video	4
Critical Applications	3
Excellent Effort	2
Best Effort	1
Background	0

Using 802.1p Priority to Provide QoS

The IBM SDN VE vSwitch provides Quality of Service functions based on the priority bits in a packet's VLAN header. (The priority bits are defined by the 802.1p standard within the IEEE 802.1Q VLAN header.) The 802.1p bits, if present in the packet, specify the priority to be given to packets during forwarding. Packets with a numerically higher (non-zero) priority are given forwarding preference over packets with lower priority value.

The IEEE 802.1p standard uses eight levels of priority (0-7). Priority 7 is assigned to highest priority network traffic, priorities 5-6 are assigned to delay-sensitive applications such as voice and video, and lower priorities are assigned to standard applications. A value of 0 (zero) indicates a "best effort" traffic prioritization, and this is the default when traffic priority has not been configured on your network. The vSwitch can filter packets based on the 802.1p values.

Figure 5. Layer 2 802.1q/802.1p VLAN tagged packet



Ingress packets receive a priority value, as follows:

- **Tagged packets**—switch reads the 802.1p priority in the VLAN tag.
- **Untagged packets**—switch tags the packet and assigns an 802.1p priority value, based on the port's default 802.1p priority.

QoS Implementation

The IBM SDN VE solution provides QoS per Virtual Network Identifier (VNID). QoS parameters are applied to traffic between two IBM SDN VE 5000V Distributed vSwitches. The vSwitch that sends the packet tags the outer header, which is later removed by the vSwitch that receives the packet.

The forwarding decision is based on the QoS parameters configured on the VNID of the port group to which the source VM is connected. Subsequently, VLAN header is inserted with default VLAN ID and with configured 802.1p values to the outer MAC header, and DSCP values are updated in outer IP header before sending the packet out on the uplink port.

When a vSwitch receives a packet, it removes the outer headers during decapsulation. The vSwitch also removes the VLAN tag from the packet before forwarding it to the destination VM.

QoS can be configured only on a DOVE connectivity group.

Rate Limiting

Rate Limiting is a mechanism to provide QoS. Outbound network traffic can be controlled by configuring a bandwidth limit. Excess traffic is held in queues and the flow of traffic thereby controlled.

Rate limiting defines the average bandwidth, peak bandwidth, and burst rate of the associated port or port group:

- **Average bandwidth:** The bandwidth in kilobytes per second (KBps) allowed for traffic on a port. The default value is 0 KBps.
- **Peak bandwidth:** A maximum amount of bandwidth (in KBps) allowed when traffic has reached the average bandwidth. The default value is 0 KBps.
- **Burst rate:** The amount of data (in Kilobytes) that is allowed to be transmitted when traffic has reached peak bandwidth rate. The default value is 0 Kilobytes.

Rate limiting can be configured on individual ports. For ports that are part of a connectivity group, rate limiting must be configured for the connectivity group, if required.

Only one rate (average bandwidth or peak bandwidth) per direction can be configured on individual ports. The burst rate is set at 0 Kilobits by default, which cannot be changed.

DOVE Connectivity Group

Rate limiting for a DOVE connectivity group can be configured with all three parameters: average bandwidth, peak bandwidth, and burst rate. This can be used for both ingress and egress traffic.

Use the following commands to configure rate limiting for a DOVE profile:

```
5000V(config)# iswitch doveprof <Profile Name>
5000V(config-dvprof)# rate-limit input <Average Rate> <Peak Rate> <Burst Rate>
5000V(config-dvprof)# rate-limit output <Average Rate> <Peak Rate> <Burst Rate>
```

When the profile is exported to the 5000V Controller, the configured parameters are passed on to the VMware vCenter to shape the traffic for the profile.

Limitations

- Only one rate can be configured per direction on standalone and on VNIC profile ports that are not part of a DOVE profile.

Chapter 15. sFlow

The IBM SDN VE 5000V Distributed vSwitch supports sFlow version 5 technology for monitoring data networks. The embedded sFlow agent can be configured to provide continuous monitoring in the form of random packet sampling and time-based sampling of statistical counters for IPv4 traffic.

The vSwitch is responsible only for forwarding sFlow information. One or more separate sFlow collectors (or analyzers) are required elsewhere on the network to interpret sFlow data.

The vSwitch provides a global sampling engine and up to 31 additional sampling engines which can be customized to monitor specific ports and/or VLANs. Each sampling engine has independent sampling rates, counter poll intervals, and can be directed to different sFlow collectors.

Further, each ESX host associated with the vSwitch vDS is an sFlow sub-agent, and each ESX host also maintains independent sample-rate counters.

Enabling sFlow

To enable the sFlow feature, use the following ISCLI configuration commands:

```
5000v(config)# sflow
5000v(config-sflow)# enable
5000v(config-sflow)# agent-ip <agent IP address>
```

The *agent IP address* represents the vSwitch to the sFlow collectors and analyzers. Set the IP address of the controller management interface. You can find this address using the `show interface ip-mgmt` command.

Although enabled, actual sampling will not occur until packet sampling or statical counters sampling are configured as shown in the following sections.

Note: Communication between the vSwitch and target sFlow collectors uses established sFlow service port 6343. sFlow operation requires that any VMware firewalls or security features permit UDP port 6343 traffic between the VMware vCenter and vSwitch controller and vDS host modules. See [“Firewall Considerations” on page 176](#) for more information.

Global Packet Sampling

When global sampling is configured, the vSwitch sFlow engine samples all packets that traverse the vSwitch.

Packets are sampled only if they successfully egress the vSwitch switching fabric, either via an access port attached to a VM or via an uplink port. Packets that are dropped by ACLs or other features will not be sampled.

When a packet sample is taken, 128 bytes are copied, UDP-encapsulated, and sent to a configured sFlow collector.

Configuring Global Packet Sampling

Global packet sampling configuration is performed by setting the sample-rate and the IP address of the sFlow collector as follows:

```
SDN-VE@SDN-VE-Controller(config-sflow)# sample-rate <packet period (1-65534)>
SDN-VE@SDN-VE-Controller(config-sflow)# collector <sFlow IP address>
```

Note: Only if sample rate is configured, the vSwitch sends sample packets.

The sFlow global sampling rate can be configured to occur once each 1 to 65534 packets. A sampling rate of 256 means that one sample will be taken for approximately every 256 packets received.

The sampling rate is statistical. It is possible to have slightly more or fewer samples sent to the collector for any specific group of packets (especially under low traffic conditions). The actual sample rate becomes most accurate over time, and under higher traffic flow.

Note: Although sFlow sampling is not generally a CPU-intensive operation, configuring extremely fast sampling rates on ports under heavy traffic loads can cause high CPU utilization on the controller or ESX hosts. Use larger rate values of 256 or more for ports that experience heavy traffic.

Disabling Global Network Sampling

To disable global packet sampling while leaving other sFlow features operational, negate the sample-rate as follows:

```
SDN-VE@SDN-VE-Controller(config-sflow)# no sample-rate
```

Network Sampling Limitations

When combined with other features, sFlow sampling the following behaviors are expected:

- Packets that are dropped by ACLs or other features will not be sampled.
- sFlow sampling will not occur on packets that are duplicated during the port mirroring process. If sFlow sampling is enabled on a port that is configured as a port monitor, the mirrored traffic will not be sampled. However, the original packet may be sampled if sFlow network sampling is enabled on its original (non-monitor) port or VLAN destination.

Statistical Counters

Note: Sample rate must already be configured.

When global counters sampling is configured, the vSwitch sends information regarding network statistical counters to an sFlow collector (or analyzer) at regular, configurable intervals.

Configuring Global Counters Sampling

Global counters sampling configuration is performed by setting the poll interval and the IP address of the sFlow collector as follows:

```
SDN-VE@SDN-VE-Controller(config-sflow)# counter-poll <interval in seconds  
(20-65534)>  
SDN-VE@SDN-VE-Controller(config-sflow)# collector <sFlow IP address>
```

Note: If the sFlow collector IP address was previously configured for packet sampling, the `collector` command can be ignored.

When the configured polling interval has elapsed, the vSwitch will report general port statistics and port Ethernet statistics to the sFlow collector. In addition, each sub-agent (vDS host module) will send its own statistical counters data. Each sub-agent maintains an independent sFlow engine. A packet traversing one ESX host will not impact the sFlow counters on another ESX host in the same vDS.

Disabling Global Counters Sampling

To disable global counters sampling while leaving other sFlow features operational, negate the polling interval as follows:

```
SDN-VE@SDN-VE-Controller(config-sflow)# no counter-poll
```

Custom Sampling Groups

The vSwitch supports up to an additional 31 sFlow sampling engines. Each can be customized to focus packet sampling and/or counters sampling on a single port or VLAN. Each sampling engine is independent of the others. Sampling engines are configured using the sFlow `group` configuration mode.

Sampling groups are numbered grouped from 1 to 31. If any port or VLAN is assigned to multiple sampling groups, the sampling group with the lowest ID number will have priority.

Configuring Sampling Groups

Sampling group configuration is performed in the sFlow Group configuration mode as follows:

```
5000v(config-sflow)# group <group number (1-31)>  
5000v(config-sflow-group)# sample-rate <packet period (1-65534)>  
5000v(config-sflow-group)# counter-poll <interval in seconds (20-65534)>  
5000v(config-sflow-group)# collector <sFlow IP address>  
5000v(config-sflow-group)# add port <port number>  
5000v(config-sflow-group)# add vlan <VLAN number>
```

The sample-rate, poll interval, and collector IP address work the same as with global sampling ([“Global Packet Sampling” on page 171](#)), but apply to only the ports and VLANs specified for the group.

To remove previously added ports or VLANs from the group, use the appropriate `del` configuration option:

```
5000v(config-sflow-group)# del port <port number>
5000v(config-sflow-group)# del vlan <VLAN number>
```

Use the `exit` command to leave the sFlow Group configuration mode.

Enabling or Disabling All Custom Sampling Groups

Custom sampling groups require the sFlow feature to be enabled (see [“Enabling sFlow” on page 171](#)). When sFlow is disabled, custom sampling groups are inactive. However, when sFlow is enabled, custom sampling groups are independent of the global sampling engine. Custom sampling groups can be used even when global packet or global counters sampling are disabled.

Enabling or Disabling Individual Groups

The collector IP address is required for sampling. To disable both packet sampling and counters sampling simultaneously, you can negate the collector IP address in the sFlow Group configuration mode:

```
5000v(config-sflow-group)# no collector
```

Enabling or Disabling Individual Group Functions

Each custom sampling group can be independently configured packet sampling, counters sampling, or both. The sample-rate is required only if packet sampling is desired. The polling interval is required only if counters sampling is desired. To disable either or both sampling functions, use the appropriate negation command in the sFlow Group configuration mode:

```
5000v(config-sflow-group)# no sample-rate
5000v(config-sflow-group)# no counter-poll
```

Order of Precedence

Each packet will be considered for sampling or counting no more than once. There is no duplication between the global sFlow engine or any of the customer sFlow sampling groups. Whether a packet is considered by the global engine or by one of the custom groups is based on the following priorities:

1. Custom group port matching.

If the packet egress port matches a port assigned to a custom sampling group, the packet will participate only in that group's sampling process.

If the egress port is included in more than one group, the group with the lowest ID is given priority. For example, group 1 has a higher priority than group 10. If the port belongs to both groups, sampling or counting of that packet will be processed according to group 1 sample-rate and counted only toward group 1 statistics.

2. Custom group VLAN matching.

Upon switch egress, if the packet's VLAN matches a VLAN assigned to a custom sampling group, the packet will participate only in that group's sampling process.

And just as with port matching, if the packet's VLAN is included in more than one custom sampling group, the group with the lowest ID is given priority.

3. Global sFlow engine

The packet will be processed for sampling and/or statistical counting only at the level where the first match is found. If a packet matches one level but is subsequently not selected to be forwarded to the sFlow collector (based on the group's sample-rate), it will not be considered for packet sampling in any other sFlow engine. Similarly, at that same match level, if the packet is not selected for statistical counting (such as when no polling interval is configured), it will not be counted at any other group or level.

sFlow Configuration Information

To obtain information about the current state of sFlow configuration, use the following global show command (shown with sample output):

```
5000v# show sflow
sFlow sampling is globally enabled
  Global sampling rate:      1 in 250 packets
  Global counter polling rate: not configured
  Collector: 172.31.46.40
  Agent IP:  172.31.38.155

Group 10: sample rate is 100, poll is 60, collector is 10.100.200.150
  ports: 10 20
  vlans: 10
5000v(config-sflow-group)# no counter-poll
```

sFlow Configuration Example

In the following example, a customer sampling group is configured. Only packets that egress port 10 or 20 (regardless of VLAN), or on VLAN 10 (regardless of port) are considered. For packets that match those criteria, 1 in 200 packets are sampled, and statistical counters are collected every two minutes.

1. Enable sFlow:

```
5000v(config)# sflow
5000v(config-sflow)# enable
```

2. Set the IP address the switch will use to identify itself to the sFlow collector:

```
5000v(config-sflow)# agent-ip 10.100.200.10
```

3. Enter the sFlow Group configuration mode:

```
5000v(config-sflow)# group 10
```

4. Set the sample-rate and polling interval

```
5000v(config-sflow-group-10)# sample-rate 200
5000v(config-sflow-group-10)# counter-poll 120
```

5. Specify the IP address of the sFlow collector or analyzer:

```
5000v(config-sflow-group-10)# collector 10.100.200.150
```

6. Add the appropriate ports and VLANs to the group:

```
5000v(config-sflow-group-10)# add port 10
5000v(config-sflow-group-10)# add port 20
5000v(config-sflow-group-10)# add vlan 10
```

7. Check the configuration information and exit the sub-modes.

```
5000v(config-sflow-group-10)# show sflow
5000v(config-sflow-group-10)# exit
5000v(config-sflow)# exit
5000v(config)#
```

Firewall Considerations

In order for the sFlow feature to work correctly, UDP traffic to destination port 6343 on the target sFlow collectors must be permitted through any VMware firewall or security features that are configured in the network.

For convenience, an example configuration is provided in the following section. However, the example may not apply in all environments. It is recommended that you defer to the appropriate VMware documented for reconfiguring security and firewall options on the ESX hypervisor.

The VMware documentation to configure the firewall can be found at the following URL:

<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-security-guide.pdf>

Example ESX hypervisor configuration

Perform the following tasks on the ESX hypervisor (login via SSH).

1. Create the file `/etc/vmware/firewall/ibm5000V.xml` with the following content:

```
<ConfigRoot>
  <service>
    <id>ibm5000VsFlow</id>
    <rule id='0000'>
      <direction>outbound</direction>
      <protocol>udp</protocol>
      <porttype>dst</porttype>
      <port>6343</port>
    </rule>
    <enabled>true</enabled>
    <required>false</required>
  </service>
</ConfigRoot>
```

2. At the ESXCLI, perform a network firewall refresh .

```
# esxcli network firewall refresh
```

3. Verify that the `ibm5000VsFlow` entry is now in the list and enabled (as shown by the final line of the sample output):

```
# esxcli network firewall ruleset list

...
remoteSerialPort      false
ibm5000VsFlow         true
```

4. Disable `allow-all` for the ruleset (unless you want to enable the firewall so that sFlow packets can go to any destination):

```
# esxcli network firewall ruleset set --allowed-all false
--ruleset-id="ibm5000VsFlow"
```

5. Enable the ruleset for the target sFlow collector's subnet or IP address. In this example, the collector is in 172.30.0.0/24:

```
# esxcli network firewall ruleset allowedip add
--ip-address=172.30.0.0/24 --ruleset-id="ibm5000VsFlow"
```

Replace 172.30.0.0/24 with an address relevant to your deployment.

Repeat this command for each additional IP address or subnet you wish to add.

6. Verify that the ruleset is properly applied:

```
# esxcli network firewall ruleset allowedip list

...
remoteSerialPort      All
ibm5000VsFlow         172.30.0.0/24
```

7. Configure sFlow on the switch as described in the other sections of this chapter.

Once full configuration is complete, sFlow packets should traverse the ESX firewall and arrive at the sFlow collector. Please consult the VMware documentation for best practices on reconfiguring that product's firewall.

Chapter 16. TCP Segmentation Offload

TCP Segmentation Offload (TSO) breaks down large groups of data (TCP packets) sent over a network into smaller segments. TSO improves network performance by reducing the CPU overhead.

The Network Interface Controller (NIC) must be capable of handling TSO. When TSO is enabled, the NIC divides large data into TCP segments. If TSO is disabled, the CPU does the segmentation.

The IBM SDN VE solution provides TSO functionality whereby data from the guest operating system sent to the IBM SDN VE 5000V Distributed vSwitch is segmented by a physical adapter capable of performing TSO.

The hosts on which the 5000V Host module is installed must have VMware vSphere Hypervisor ESXi 5.5. You must connect these hosts to the vSwitch.

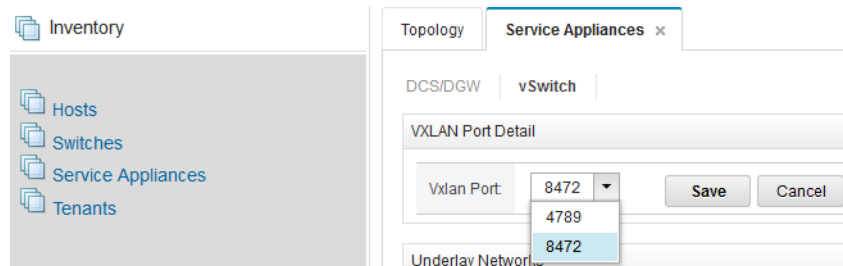
All uplink ports that are part of a port group must have the same TSO setting: enabled or disabled. If any port has a different setting, TSO is automatically disabled and the CPU performs the packet segmentation.

VXLAN Port

Virtual eXtensible Local Area Network (VXLAN) is a tunneling mechanism to overlay Layer 2 networks on top of Layer 3 networks. You can set the UDP port to 4789 or 8472. Default port is 8472.

Set the UDP port as follows:

1. Login to the controller GUI: <https://<Controller HA external IPv4 address>:8443>.
2. Select **Inventory** > **Service Appliances** > **vSwitch**.
3. Select **Edit** under the **VXLAN Port Detail** section.



4. Select the port.
5. Select **Save**.

Chapter 17. Virtual Router Redundancy Protocol

The IBM SDN VE solution supports IPv4 high-availability (HA) network topologies through implementation of the Virtual Router Redundancy Protocol (VRRP).

VRRP Overview

VRRP enables redundant router configurations within a LAN, providing alternate router paths for a host to eliminate single points-of-failure within a network. Each participating VRRP-capable routing device is configured with the same virtual router IPv4 address and ID number. One of the virtual routers is elected as the master, based on a number of priority criteria, and assumes control of the shared virtual router IPv4 address. If the master fails, one of the backup virtual routers will take control of the virtual router IPv4 address and actively process traffic addressed to it.

With VRRP, Virtual Interface Routers (VIR) allow two VRRP routers to share an IP interface across the routers. VIRs provide a single Destination IPv4 (DIP) address for upstream routers to reach various servers, and provide a virtual default Gateway for the servers.

VRRP Components

Each physical router running VRRP is known as a *VRRP router*.

Virtual Router

Two or more VRRP routers can be configured to form a *virtual router* (RFC 2338). Each VRRP router may participate in one or more virtual routers. Each virtual router consists of a user-configured *virtual router identifier* (VRID) and an IPv4 address.

Master and Backup Virtual Router

Within each virtual router, one VRRP router is selected to be the virtual router master. See [“Selecting the Master VRRP Router” on page 181](#) for an explanation of the selection process.

Note: If the IPv4 address owner is available, it will always become the virtual router master.

The virtual router master forwards packets sent to the virtual router. It also responds to Address Resolution Protocol (ARP) requests sent to the virtual router's IPv4 address. Finally, the virtual router master sends out periodic advertisements to let other VRRP routers know it is alive and its priority.

Within a virtual router, the VRRP routers not selected to be the master are known as virtual router backups. If the virtual router master fails, one of the virtual router backups becomes the master and assumes its responsibilities.

Selecting the Master VRRP Router

Each VRRP router is configured with a priority between 1–254. A bidding process determines which VRRP router is or becomes the master—the VRRP router with the highest priority.

The master periodically sends advertisements to an IPv4 multicast address. As long as the backups receive these advertisements, they remain in the backup state. If a backup does not receive an advertisement for three advertisement intervals, it

initiates a bidding process to determine which VRRP router has the highest priority and takes over as master. In addition to the three advertisement intervals, a manually set holdoff time can further delay the backups from assuming the master status.

A backup router can stop receiving advertisements for one of two reasons—the master can be down, or all communications links between the master and the backup can be down. If the master has failed, it is clearly desirable for the backup (or one of the backups, if there is more than one) to become the master.

VRRP Implementation

IBM SDN VE solution supports active-standby configuration.

In an IBM SDN VE setup, Distributed Gateways (DGWs) provide Source Network Address Translation (SNAT) and forwarding services to connected Virtual Network Identifiers (VNIDs). Each DGW can provide services to one or more VNIDs. When two DGWs are configured for HA, one acts as a Master and the other as a Backup. Together, they provide service to all the connected VNIDs. A DGW can act as a backup for only one Master DGW.

DGWs configured for HA use a virtual tunnel endpoint IP (Virtual TEP) address to receive packets from DOVE virtual switches.

Both the DGWs and associated VNIDs must be in the same IP subnetwork. The networks to which the DGWs connect to must support multicast.

Configuring VRRP

The DOVE Connectivity Gateways can be configured using the controller GUI as follows.

1. Login to the controller GUI: `https://<Controller HA external IPv4 address>:8443`.
2. Select the tenant.
3. Select **Inventory > Service Appliances > DCS/DGW**.
4. Select **Config VRRP** from the **DGW List** section. The VRRP List window is displayed.

5. Select **Create VRRP** (📄) icon from the top right corner.

Create VRRP

* Virtual TEP:

* Virtual Ext IP:

* GW 1 Index:

* GW 1 Priority:

* GW 2 Index:

* GW 2 Priority:

* Virtual Router ID:

OK

Cancel

Specify the following:

Table 22. VRRP Specifications

Component	Description
Virtual TEP	Virtual IP address of tunnel endpoint.
Virtual Ext IP	External virtual IP address.
GW 1 Index	Index of the DGW appliance that you want to configure as master.
GW 1 Priority	<p>Priority of the master DGW appliance. Specify a value between 1 and 254. The DGW with a higher priority is selected as the Master.</p> <p>Note: In this release, the priority value configured on a DGW is not effective. DGW with the smaller IP address is always selected as the Master.</p>
GW 2 Index	Index of the DGW appliance that you want to configure as backup.
GW 2 Priority	<p>Priority of the backup DGW appliance. Specify a value between 1 and 254. The DGW with a higher priority is selected as the Master.</p> <p>Note: In this release, the priority value configured on a DGW is not effective. DGW with the smaller IP address is always selected as the Master.</p>
Virtual Router ID	Specify a Virtual Router ID. Both DGWs are configured with this VRID.

6. Select **OK**.

Part 3: Command Reference

Chapter 18. Command Basics

The IBM SDN VE system is ready to perform basic networking functions after initial installation. Some of the more advanced features, however, require some administrative configuration before they can be used effectively.

The SDN VE system provides a Command-Line Interface (CLI). Using a basic terminal, the CLI allows you to view information and statistics about the virtual network, and to perform any necessary configuration.

The CLI is available on any installed controller where primary information and configuration is performed. A more limited CLI is also available on any installed Distributed Service Appliance (DSA), and is used mainly for initial setup purposes.

This chapter explains how to use the CLI available in the controller and DSA modules.

Login

CLI access is controlled through the use of a login name and password. Once you are connected to the system via SSH, you are prompted to enter a login name and password.

Default user name: **admin**

Default password: **admin**

Note: It is recommended that you change all default system password after initial configuration and as regularly as required under your network security policies.

Command Modes

Once logged in to the Unified Controller, the CLI commands are organized by context. The various contexts, or *modes*, are organized in hierarchical fashion. The modes, their identifying prompts, and their navigational commands are shown in the following figure:

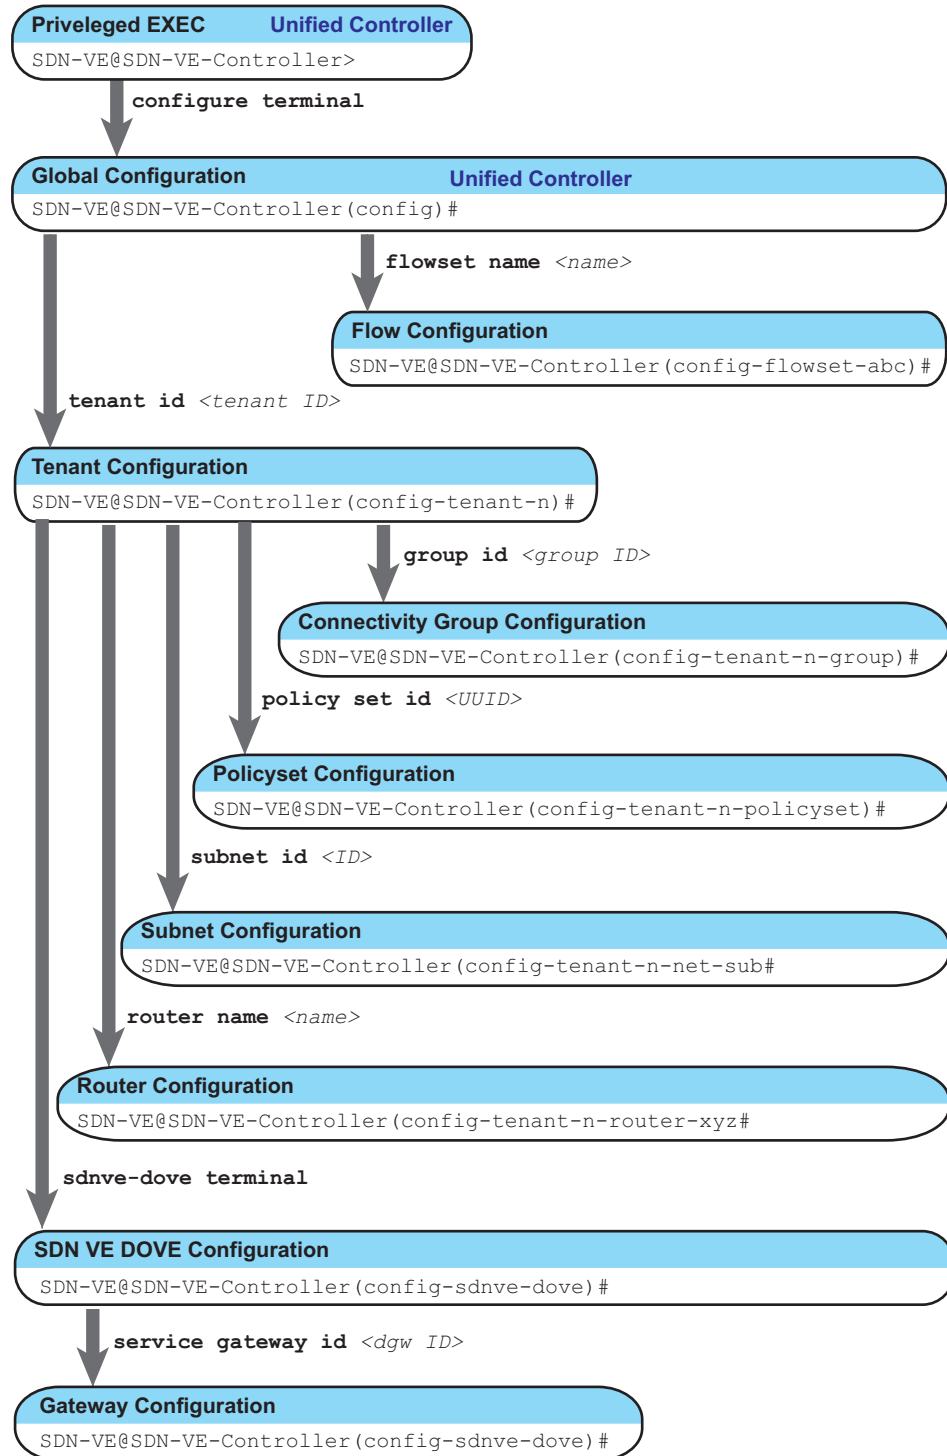


Figure 6. CLI Command Modes

Unified Controller

Privilege EXEC Mode

This is the initial access mode granted upon login. This mode is used to collect information and execute limited operational commands. To avoid accidental configuration changes, commands that affect the permanent configuration are not permitted in this mode. This mode is available in Unified Controller node.

Identifying prompt:

```
SDN-VE@SDN-VE-Controller>
```

Mode navigation commands:

- `configure terminal`
Enter Configuration mode.
- `exit` or `quit`
Quit the CLI session.

Global Configuration Mode

This mode allows you to make changes to the running configuration. All changes take effect immediately (unless otherwise noted) and survive a reset of the system. This mode is available in the Unified Controller node.

Use the following Privileged EXEC mode command to access the Global Configuration mode:

```
SDN-VE@SDN-VE-Controller# configure terminal
```

Identifying prompt:

```
SDN-VE@SDN-VE-Controller(config)#
```

Several sub-modes are available from the Global Configuration mode. Each mode provides a specific set of commands.

Mode navigation commands:

- `flowset name <name>`
Enter the Flow Group Configuration mode. This mode is used for configuring flows.
- `tenant id <tenant ID>`
Enter the Tenant Configuration mode. Tenant-specific information is configured in this mode.
- `sdnve-dove terminal`
Enter the SDN VE DOVE Configuration mode.
- `exit`
Return to Privileged EXEC mode.
- `quit`
Quit the CLI session.

SDN VE DOVE Controller

Enter this mode from the Unified Controller CLI. This mode is used for controller key configurations such as external IP, High Availability, Gateway, underlay networks.

SDN VE DOVE Configuration Mode

Use the following command to access the SDN VE DOVE Configuration mode:

```
SDN-VE@SDN-VE-Controller(config)# sdnve-dove terminal
```

Identifying prompt:

```
SDN-VE@SDN-VE-Controller(config-sdnve-dove)#
```

Gateway Configuration

Define forwarding rules for gateway. Gateways connect the SDN VE virtual network with traditional (non-virtual) networks.

```
SDN-VE@SDN-VE-Controller(config-sdnve-dove)#  
service gateway id <gateway ID> ?
```

where <ID> identifies the gateway you wish to configure.

Mode navigation commands:

- `exit`
Return to Global Configuration mode.
- `quit`
Quit the CLI session.

Global Commands

Some basic commands are recognized throughout all CLI command modes. These commands are useful for navigating through the interface.

Table 23. Description of Global Commands

Command	Action
<code>?</code>	List the commands available in the current mode, or when placed after a command keyword, provide further information about command options.
<code>copy</code>	Save the current configuration. <code>copy running-config startup-config</code>
<code>exit</code>	Go up one level in the command mode structure. If already at the top level, exit from the command line interface and log out.
<code>history</code>	Displays the commands executed in the current session. <code>history [<list size>]</code>
<code>no</code>	Use the <code>no</code> form of a command to delete a configuration. Applicable only to Global Configuration mode or higher. <code>no external-ip</code>
<code>quit</code>	Exit from the CLI and log out.
<code>show</code>	shows configuration information. See Chapter 19, "Show Commands" .

Idle Timeout

By default, the CLI session will be disconnected after ten minutes of inactivity. This function is controlled by the following command, which can be set from 1 to 30 minutes:

```
SDN-VE-Controller(config)# system cli timeout <1-30>
```

Command mode:

Global Configuration mode or higher.

Chapter 19. Show Commands

Once you have logged in to a controller, you can view system configuration and statistical information using a variety of CLI `show` commands. The `show` commands are restricted from the User EXEC mode, but most are available globally in all other command modes.

Please note that the output may exceed 10K of data, depending on your configuration.

If you want to capture the data to a file, such as for support or diagnostic purposes, set the communication software on your workstation to capture session data prior to issuing the command.

The remainder of this chapter discusses how to use each of the information-specific CLI `show` commands.

Cluster Information

`show cluster info`

Shows information about the nodes in a cluster.

Syntax:

```
show cluster info
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show cluster info

Cluster name : test
Node          Role          Status      Sync
-----
9.121.62.116  standby    online      completed
9.121.62.118* active      online      completed
```

SPARTA Information

`show connectivity path`

Shows the SPARTA connectivity path between two hosts specified by source and destination MAC addresses.

Syntax:

```
show connectivity path src-mac <Mac address of source host> dest-mac
<MAC address of destination host>
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show connectivity path src-mac 00:00:00:00:00:01
dest-mac 00:00:00:00:00:02
=====
* 00:00:00:00:00:00:00:01
+ (2) -- (2) 00:00:00:00:00:00:00:02
=====
```

show connectivity tree

Shows the SPARTA connectivity tree for the specific destination MAC address.

Syntax:

show connectivity tree dest-mac *<MAC address of destination host>*

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show connectivity tree dest-mac 00:00:00:00:00:03
=====
* 00:00:00:00:00:00:00:03
+ (2) -- (3) 00:00:00:00:00:00:00:02
+ (2) -- (2) 00:00:00:00:00:00:00:01
=====
```

Flow Information

show flow

Shows information about the static and dynamic flows.

Syntax:

show flow {all|dynamic|static} switch *<Switch ID>* [count]

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show flow all switch 00:00:00:01:03:04:05:06
```

```
1. name           : **ArpPunt
   id             : 1234.5678.90123456
   base           : Yes
   priority        : 500
   installInHw     : true
   idleTimeout     : 120
   hardTimeout     : 240
   In Port         : *
   DL Src          : *
   DL Dst          : *
   DL Type         : *
   DL Vlan         : *
   Vlan-priority   : *
   NW Src          : *
   NW Dst          : *
   NW Protocol     : *
   Nw-tos          : *
   TP Src          : *
   TP Dst          : *
   Actions         : DROP
O/p Port(s)       : 1800, 8909...
O/p Vlan          : *
Set Dl Src        : *
Set Dl Dst        : *
Set Valn Pcp      : *
Set Nw Tos        : *
```

```
2. name           : tstFlow
   id             : 1234.5678.90123478
   base           : No
   priority        : 700
   installInHw     : true
   idleTimeout     : 180
   hardTimeout     : 360
   In Port         : 4
   DL Src          : *
   DL Dst          : *
   DL Type         : *
   DL Vlan         : *
   NW Src          : *
   NW Dst          : *
   NW Protocol     : *
   Nw-tos          : *
   TP Src          : *
   TP Dst          : *
   Actions         : OUTPUT
O/p Port(s)       : 6, 7
O/p Vlan          : *
Set Dl Src        : *
Set Dl Dst        : *
Set Valn Pcp      : *
Set Nw Tos        : *
```

```
SDN-VE @SDN-VE-Controller > show flow all switch 00:00:00:01:03:04:05:06 count
```

```
Number of flows = 8
```

show flowset

Shows information about flow groups.

Syntax:

```
show flowset {all|dynamic [name <flow group name>]  
|static [name <flow group name>]} [count]
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show flowset all  
  
1  name                : tstFlowGroup  
   priority             : 700  
   installInHw         : true  
   idleTimeout         : 180  
   hardTimeout         : 360  
   Owner                : Sparta  
Flows :  
1.  name                : tstFlowGroup-flow1  
    id                  : 1234.5678.90343456  
    node                : node1  
    priority            : 700  
    installInHw         : true  
    idleTimeout         : 180  
    hardTimeout         : 360  
    In Port             : 2  
    DL Src              : *  
    DL Dst              : *  
    DL Type             : *  
    DL Vlan             : *  
    Vlan-priority       : *  
    NW Src              : *  
    NW Dst              : *  
    NW Protocol         : *  
    Nw-tos              : *  
    TP Src              : *  
    TP Dst              : *  
    Actions             : OUTPUT  
    O/p Port(s)         : 6  
    O/p Vlan            : *  
    Set Dl Src          : *  
    Set Dl Dst          : *  
    Set Valn Pcp        : *  
    Set Nw Tos          : *  
    Set Dl Dst          : *  
    Set Valn Pcp        : *  
    Set Nw Tos          : *  
(cont.)...
```

```

...(cont.)
2 name                : tstFlowGroup2
  priority             : 800
  installInHw         : true
  idleTimeout         : 120
  hardTimeout         : 240
  Owner               : L3
Flows :
1. name                : tstFlowGroup2-flow1
   id                  : 1234.5678.90323456
   node                : node3
   priority            : 700
   installInHw         : true
   idleTimeout         : 180
   hardTimeout         : 360
   In Port             : 2
   DL Src              : *
   DL Dst              : *
   DL Type             : *
   DL Vlan             : *
   Vlan-priority       : *
   NW Src              : *
   NW Dst              : *
   NW Protocol         : *
   Nw-tos              : *
   TP Src              : *
   TP Dst              : *
   Actions             : OUTPUT
   O/p Port(s)         : 9
   O/p Vlan            : *
   Set Dl Src          : *
                       Set Dl Dst          : *
                       Set Valn Pcp        : *
   Set Nw Tos          : *

```

```

SDN-VE @SDN-VE-Controller > show flowset dynamic <flow group name>

```

show statistics flow

Shows statistical information about all flows or a particular flow.

Syntax:

```

show statistics flow [switch <switch ID> | flow <flow ID> ][flow <flow
ID> | switch <switch ID> ]

```

Command mode:

Privileged Executive

Example:

```
SDN-VE @SDN-VE-Controller > show statistics flow switch 00:00:00:01:03:04:05:06 id
050b4f82-726b-3215-b3fe-40795ded5440

flow id : 050b4f82-726b-3215-b3fe-40795ded5440
table-id      : 1
duration      : 25 secs
priority      : 100
idle-timeout  : 0
hard-timeout  : 0
packets       : 12345
bytes         : 1234500
```

Connectivity Group Information

show group

Show information about the connectivity groups. Information can be viewed based on the connectivity group ID or name. All connectivity groups configured for a tenant can also be viewed.

Syntax:

```
show group [id | name <Connectivity group name> | tenant id <tenant ID>]
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show group

Id                : 10
Name              : SDN9990104_CG1
Vnid             : 0
Tenant_id        : 1000
Subnets         : 100, 102
Admin State      : true
Status           : ACTIVE
Group_type       : dedicated
isNeutron        : false
Waypoint         : false

Id                : 96a5e853-f5dd-38b7-9c65-431d3b8888da
Name              : txpnt1
Vnid             : 0
Tenant_id        : 1000
Admin State      : true
Status           : ACTIVE
Group_type       : dedicated
isNeutron        : false
Waypoint         : true

Id                : a1568745-50f2-3c66-bb8f-96cc7bfc4500
Name              : rnat2
Vnid             : 0
Tenant_id        : 1000
Subnets         : 01bb28d5-4772-3cba-8a0c-df1f5920b1bf,
2799d4ff-ebcb-3c7b-830f-1ad6ceb620fd
Admin State      : true
Status           : ACTIVE
Group_type       : dedicated
isNeutron        : false
Waypoint         : true

Id                : f1a9a170-7623-35e7-8a15-85bfc183b5bd
Name              : rtd3
Vnid             : 0
Tenant_id        : 1000
Subnets         : 2799d4ff-ebcb-3c7b-830f-1ad6ceb620fd
Admin State      : true
Status           : ACTIVE
Group_type       : dedicated
isNeutron        : false
Waypoint         : true

Id                : f5c400ee-38e4-3e4a-8b56-5e4398c4b2a1
Name              : DT
Vnid             : 4
Tenant_id        : 100
Admin State      : true
Status           : ACTIVE
Group_type       : dedicated
isNeutron        : false
Waypoint         : true

(END)
```

```
SDN-VE @SDN-VE-Controller > show group tenant id 1000
```

```
Id                : 10
Name              : SDN9990104_CG1
Vnid             : 0
Tenant_id        : 1000
Subnets         : 100, 102
Admin State      : true
Status           : ACTIVE
Group_type       : dedicated
isNeutron        : false
Waypoint         : false
```

```
Id                : 96a5e853-f5dd-38b7-9c65-431d3b8888da
Name              : txpnt1
Vnid             : 0
Tenant_id        : 1000
Admin State      : true
Status           : ACTIVE
Group_type       : dedicated
isNeutron        : false
Waypoint         : true
```

```
Id                : a1568745-50f2-3c66-bb8f-96cc7bfc4500
Name              : rnat2
Vnid             : 0
Tenant_id        : 1000
Subnets         : 01bb28d5-4772-3cba-8a0c-df1f5920b1bf,
2799d4ff-ebcb-3c7b-830f-1ad6ceb620fd
Admin State      : true
Status           : ACTIVE
Group_type       : dedicated
isNeutron        : false
Waypoint         : true
```

```
Id                : f1a9a170-7623-35e7-8a15-85bfc183b5bd
Name              : rtd3
Vnid             : 0
Tenant_id        : 1000
Subnets         : 2799d4ff-ebcb-3c7b-830f-1ad6ceb620fd
Admin State      : true
Status           : ACTIVE
Group_type       : dedicated
isNeutron        : false
Waypoint         : true
```

```
(END)
```

Host Information

show host

Shows information about hosts. Information about a particular host based on the IP address, or information about all active and inactive hosts can be viewed.

Syntax:

```
show host inactive all
show host active {all | ip <Host IP address>}
```

Command mode:

Privileged Executive

Example:

SDN-VE @SDN-VE-Controller > show host active all					
MAC Address	IP Address	VLAN	User Configured	Switch Id	Port Id
00:00:00:00:00:02	10.0.0.2	0	false	00:00:00:00:00:00:00:02	2
00:00:00:00:00:04	10.0.0.4	0	false	00:00:00:00:00:00:00:03	2
00:00:00:00:00:03	10.0.0.3	0	false	00:00:00:00:00:00:00:03	1
00:00:00:00:00:01	10.0.0.1	0	false	00:00:00:00:00:00:00:02	1

LDAP Server Information

show ldap

Shows information about LDAP servers.

Syntax:

```
show ldap
```

Command mode:

All

Example:

SDN-VE @SDN-VE-Controller > show ldap	
ldap	: enabled
server	: 9.1.2.5
domain	: ou=people,dc=mydomain,dc=com

Log Information

Log Levels

show level

Shows the log levels configured for various services/events.

Syntax:

```
show level {all | <service/event name>}
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show level all
```

LogName	Level
ping	INFO
protocol_plugin_openflow	INFO
infinispan	ERROR
loggingBridge	WARN
logical_groups	INFO
web	INFO
nist	INFO
lnp_topology	INFO
security	INFO
dove_config	
clustering	INFO
license	INFO
sal	INFO
smarttime	INFO
hosttracker	INFO
broadcast	INFO
arphandler	INFO
commons	INFO
script_interface	INFO
ofp_processor	INFO
flow_reaper	INFO
pipeline	INFO
layer3	INFO
layer2	INFO
policymanager	INFO
waypoint	DEBUG
proxy	INFO
odl_services	INFO
restore	INFO
flowgroupsmanager	
interface_manager	INFO
replication	INFO
topology	INFO
multicast	INFO
usermanager	INFO

```
(END)
```

```
SDN-VE @SDN-VE-Controller > show level waypoint  
Logger level : DEBUG
```

View Logs

show log

Shows the logged information.

Syntax:

`show log [<number of lines to display>]`

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show log 10  
  
2014-07-09 06:57:18.493 GMT [Hashed wheel timer #1] WARN  
o.o.c.p.o.c.i.TrafficStatisticsHandler - ERROR_MSG_RECEIVED 0  
2014-07-09 06:57:18.493 GMT [Hashed wheel timer #1] WARN  
o.o.c.p.o.c.i.TrafficStatisticsHandler - MESSAGE_RECEIVED 0  
2014-07-09 06:57:18.493 GMT [Hashed wheel timer #1] WARN  
o.o.c.p.o.c.i.TrafficStatisticsHandler -  
CONTROLLER_INITIATED_SWITCH_DISCONNECTS_MSGPARSE_EXCEPTION 0  
2014-07-09 06:57:18.493 GMT [Hashed wheel timer #1] WARN  
o.o.c.p.o.c.i.TrafficStatisticsHandler - CHANNEL_WRITE_COMPLETED_EVENT 0  
2014-07-09 06:57:18.493 GMT [Hashed wheel timer #1] WARN  
o.o.c.p.o.c.i.TrafficStatisticsHandler -  
PACKET_OUT_DROPPED_WRITE_OVERLOAD 0  
2014-07-09 06:57:18.493 GMT [Hashed wheel timer #1] WARN  
o.o.c.p.o.c.i.TrafficStatisticsHandler -  
CONTROLLER_INITIATED_SWITCH_DISCONNECTS_NO_ECHO_RESPONSE 0  
2014-07-09 06:57:18.494 GMT [Hashed wheel timer #1] WARN  
o.o.c.p.o.c.i.TrafficStatisticsHandler - SWITCH_ADDED_TO_DPID_MAP 0  
2014-07-09 06:57:18.494 GMT [Hashed wheel timer #1] WARN  
o.o.c.p.o.c.i.TrafficStatisticsHandler - HELLO_RECEIVED 0  
2014-07-09 06:57:18.494 GMT [Hashed wheel timer #1] WARN  
o.o.c.p.o.c.i.TrafficStatisticsHandler -  
DROPPING_HIGH_PRIORITY_MESSAGE_ON_WRITE_OVERLOAD 0  
2014-07-09 06:57:18.494 GMT [Hashed wheel timer #1] WARN  
o.o.c.p.o.c.i.TrafficStatisticsHandler - >>>>>Raw Counter values at  
controller END <<<<<<<
```

Multicast Information

show multicast

Shows information about the multicast groups, VLANs, and query intervals.

Syntax:

```
show multicast group {all | ip <Multicast IP address>}
show multicast query-interval
show multicast vlan
```

Command mode:

Privileged Executive; Global Configuration

Example:

```
SDN-VE @SDN-VE-Controller > show multicast group all

Multicast IP : 225.0.0.125
Tenant ID : 2
Multicast MAC : 01:00:5e:00:00:7d
FDB based : false
VLAN ID : 1
Joins : 10
Leaves : 4
Senders : 72:1c:89:1c:c3:a3
Members : de:7f:4a:05:9d:6c/10.0.0.7
           72:1c:89:1c:c3:a3/10.0.0.1 62:08:36:47:82:68/10.0.0.17
           9e:01:0d:9f:65:96/10.0.0.15 d6:fa:6d:a0:9e:09/10.0.0.21
           2a:23:78:b7:cf:7c/10.0.0.2

Multicast IP : 225.0.0.126
Tenant ID : 2
Multicast MAC : 01:00:5e:00:00:7e
FDB based : false
VLAN ID : 1
Joins : 10
Leaves : 4
Senders : 72:1c:89:1c:c3:a3
Members : de:7f:4a:05:9d:6c/10.0.0.7
           72:1c:89:1c:c3:a3/10.0.0.1 62:08:36:47:82:68/10.0.0.17
           9e:01:0d:9f:65:96/10.0.0.15 d6:fa:6d:a0:9e:09/10.0.0.21
           2a:23:78:b7:cf:7c/10.0.0.2
```

OpenFlow Information

show ofversion

Shows the OpenFlow version of the flow group manager.

Syntax:

`show ofversion`

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show ofversion  
OFVersion=1.0
```

show statistics openflow

Shows OpenFlow statistical information.

Syntax:

`show statistics openflow switch <switch ID>`

Command mode:

Privileged Executive

Example:

```
SDN-VE @SDN-VE-Controller > show statistics openflow switch 00:00:00:01:03:04:05:06

Hello Sent : 0                      Hello Received : 0
Echo-Request Sent : 0              Echo-Request Received : 0
Echo-Reply Sent : 0                Echo-Reply Received : 0
Barrier-Request : 0                Barrier-Reply : 0

Packet-In : 0
Packet-Out : 0

Flow-Removed : 0
Flow-Mod : Add : 0                Modify : 0                Delete: 0
                                      Modify-Strict : 0
Delet-Strict: 0

Port-Status : Add : 0            Delete : 0            Modify : 0

Stats-Request : Desc : 0          Flow : 0          Aggregate : 0
                                      Table : 0          Port : 0
Stats-Reply : Desc : 0          Flow : 0          Aggregate : 0
                                      Table : 0          Port : 0

Hello Failed Sent : 0              Hello Failed Recv : 0

Bad Requests : Version : 0        Type : 0        Stat : 0        Vendor : 0
Subtype : 0
                                      Len : 0        Buffer-empty : 0
Buffer-Unknown : 0

Flow Mod Failed : Table-Full : 0    Overlap : 0    Permission Error : 0
                                      Bad-Command ; 0    Unsupported : 0

Bad Action : Bad-Type : 0        Bad-Len : 0        Bad-Out-Port : 0
Bad-Argument : 0 Too-many : 0
```

Port Information

show port

Shows information about individual ports and ports that belong to a connectivity group.

Syntax:

```
show port [group id <group ID> | port id <port ID> | name <port name> ]
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show port group id 2
ID : 1
Group Id : 2
Name : port1
Admin State : false
Status : BUILD
Mac Address : 00:00:00:00:00:01
IP Address :
Tenant Id : 2
```

show statistics port

Shows statistical information about all ports or a particular port.

Syntax:

```
show statistics port [port <port ID> | switch <switch ID> ][switch
<switch ID> | port <port ID> ]
```

Command mode:

Privileged Executive

Example:

```
SDN-VE @SDN-VE-Controller > show statistics port switch 00:00:01:02:03:04:05:06 port
10
switch 00:00:01:02:03:04:05:06    port 10
  Rx-packets : 8876541234    Tx-packets :
  Rx-bytes    :              Tx-bytes    :
  Rx-dropped:              Tx-dropped:
  Rx-errors   :              Tx-errors   :
  Rx-frame-error:          collisions   :
  Rx-overflow :
  Rx-crc-error:              collisions   :
```

RADIUS Server Information

show radius

Shows information about the RADIUS server.

Syntax:

```
show radius
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show radius
radius      : enabled
server     : 9.1.2.5
key        : 9876345321
```

Replication Session Information

show replication

Shows information about all replication sessions.

Syntax:

```
show replication
```

Command mode:

Privileged Executive; Tenant Configuration

Example:

```
SDN-VE @SDN-VE-Controller > show replication

session-name      : xyz
src-mac/src-ip    : 10.1.1.2
src-port          : *
src-protocol      : tcp
dest-mac/dest-ip  : *
dest-port         : 22
src-dest-tenant   : 1234
target-mac/target-ip : 2a:2b:2c:01:2c:04
target-tenant     : 1234
state             : started
Mode              : Replicate

session-name      : mon-1
src-mac/src-ip    : 10.1.1.1
src-port          : *
src-protocol      : tcp
dest-mac/dest-ip  : *
dest-port         : 22
src-dest-tenant   : 1234
target-mac/target-ip : 2a:2b:2c:01:2c:05
target-tenant     : 1234
state             : started
Mode              : Replicate
```

Subnet Information

show subnet

Shows information about subnets in a connectivity group or tenant. Information about all subnets or a particular subnet can be viewed based on the subnet ID or subnet name.

Syntax:

```
show subnet [group id <group ID>] [id <subnet ID>] [name <subnet name>]  
[tenant id <tenant ID>]
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show subnet id 1  
  
Id : 1  
Tenant Id : 2  
Name : sub1  
CIDR : 10.0.0.1/24  
Gateway Ip : 10.0.0.1  
IP Version : 4
```

Switch Information

show switch

Shows information about all switches or a particular switch that are/is registered with the controller.

Syntax:

```
show switch [all] [dpId <switch ID>] [name <switch name>]
```

Command mode:

Privileged Executive

Example:

```
SDN-VE @SDN-VE-Controller > show switch all  
  
Dpid                               Name      Type      MAC Address      Tier  
-----  
00:00:00:00:00:00:00:01           OF        00:00:00:00:00:01      1  
0:00:00:00:00:00:00:02           OF        00:00:00:00:00:01      1
```

```
SDN-VE @SDN-VE-Controller > show switch dpId 00:00:00:00:00:00:03

Node Properties :
dpId      :00:00:00:00:00:00:03
Type      :OF

Node Connector Properties :-
-----
Name      Node Connector Id Node Connector Type Node Connector Latency Node Connector
Bandwidth
-----
s3-eth1   1                  OF                  10000000000
s3        0                  SW
s3-eth3   3                  OF                  10000000000
s3-eth2   2                  OF                  10000000000
```

show statistics switch

Shows statistical information about all switches or a particular switch.

Syntax:

```
show statistics switch [<switch ID>]
```

Command mode:

Privileged Executive

Example:

```
SDN-VE @SDN-VE-Controller > show statistics switch 00:00:00:00:00:00:01

Switch : 00:00:00:00:00:00:01
  Rx-packets : 8876541234    Tx-packets :
  Rx-bytes   :                Tx-bytes   :
  Rx-dropped:                Tx-dropped:
  Rx-errors  :                Tx-errors  :
  Rx-frame-error:
  Rx-overflow:
  Rx-crc-error:                collisions :
```

System Information

show system acknowledgment

Shows information about the General Public License.

Syntax:

```
show system acknowledgment
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show system acknowledgment
```

```
Licenses and Attributions Document
```

```
Created: Thu Jan 14th 2014
```

```
=====
```

```
.  
.   
.
```

show system cpu

Shows information about the CPU.

Syntax:

```
show system cpu info
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show system cpu info
```

```
Processor      : 0  
Vendor         : GenuineIntel  
CPU family    : 6  
Model         : 30  
Model name    : Intel(R) Xeon(R) CPU X3460 @ 2.80GHz  
CPU MHz       : 2792.984  
CPU Cache size : 8192 KB
```

```
Processor      : 1  
Vendor         : GenuineIntel  
CPU family    : 6  
Model         : 30  
Model name    : Intel(R) Xeon(R) CPU X3460 @ 2.80GHz  
CPU MHz       : 2792.984  
CPU Cache size : 8192 KB
```

show system disk

Shows information about the disk usage.

Syntax:

```
show system disk info
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show system disk info
```

Mount-point	1K Blocks	Available	% Used
/	1032088	410340	60%
/flash	16513960	16101180	2%

show system ipmgmt

Shows IP, name server, and nexthop information about the management interfaces. Information about all management interfaces or a particular management interface can be viewed.

Syntax:

```
show system ipmgmt {ip | nameserver | nexthop}
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show system ipmgmt ip
```

Mode	: STATIC
IP	: 9.121.62.118
Mask	: 255.255.254.0
Nexthop	: 9.121.62.1
DNS	: *

```
SDN-VE @SDN-VE-Controller > show system ipmgmt nameserver
```

NAMESERVER INFORMATION	
ID	IP
1	10.44.11.22

```
SDN-VE @SDN-VE-Controller > show system ipmgmt nexthop
```

Nexthop	: 9.121.62.1
---------	--------------

show system license

Shows license information.

Syntax:

```
show system license
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show system license

License-1
-----
Version   :1
Feature   :OF
Validity  :Valid - Permanent
Expiry    :None

License-2
-----
Version   :1
Feature   :SDN_VE_xKVM
Validity  :Valid - Permanent
Expiry    :None
```

show system memory

Shows system memory information.

Syntax:

```
show system memory info
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show system memory info

Total Memory:                8184992 kB
Free Memory:                  5382544 kB
```

show system network

Shows information about port listeners and network statistics.

Syntax:

```
show system network {port listeners | statistics}
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show system network port listeners
```

Prot	Local Address	Remote Address	State	PID
tcp	127.0.0.1:9100	0.0.0.0:*	LISTEN	261
tcp	127.0.0.1:9200	0.0.0.0:*	LISTEN	262
tcp	0.0.0.0:80	0.0.0.0:*	LISTEN	202
tcp	127.0.0.1:9600	0.0.0.0:*	LISTEN	267
tcp	127.0.0.1:9000	0.0.0.0:*	LISTEN	178
tcp	9.121.62.118:40388	9.121.62.116:1964	ESTABLISHED	209
tcp	127.0.0.1:19906	127.0.0.1:1964	ESTABLISHED	269
tcp	127.0.0.1:19908	127.0.0.1:1964	ESTABLISHED	272
tcp	127.0.0.1:19907	127.0.0.1:1964	ESTABLISHED	271
tcp	9.121.62.240:22	9.79.195.189:53588	ESTABLISHED	
24548				
tcp	:::1964	:::*	LISTEN	209
tcp	::ffff:127.0.0.1:8080	:::*	LISTEN	295
tcp	::ffff:9.121.62.118:32754	:::*	LISTEN	295
tcp	:::61619	:::*	LISTEN	295
tcp	:::22	:::*	LISTEN	
1157				
tcp	::ffff:9.121.62.118:7800	:::*	LISTEN	295
tcp	::ffff:9.121.62.118:7801	:::*	LISTEN	295
tcp	:::8443	:::*	LISTEN	295
tcp	::ffff:127.0.0.1:35934	:::*	LISTEN	295
tcp	:::4126	:::*	LISTEN	209
tcp	::ffff:9.121.62.118:46943	:::*	LISTEN	295
tcp	::ffff:127.0.0.1:42977	:::*	LISTEN	295
tcp	:::12001	:::*	LISTEN	295
tcp	:::16324	:::*	LISTEN	295
tcp	::ffff:127.0.0.1:1964	::ffff:127.0.0.1:19906	ESTABLISHED	209

```
SDN-VE @SDN-VE-Controller > show system network statistics
```

```
Ip-statistics:
    Forwarding                2
    InputReceives             86755105
    InputHdrErrors            0
    InputAddrErrors           205663
    ForwardDatagrams          0
    InputDiscards             0
    OutputRequests            79027139
    OutputDiscards            0
    OutputNoRoutes            2
Icmp-statistics:
    InputMsgs                 62
    InputErrors               7
    InputEchos                0
    InputEchoReplies          0
    OutputMsgs                128
    OutputErrors              0
    OutputDestUnreaches       128
    OutputTimeExceeds         0
    OutputRedirects           0
    OutputEchos               0
    OutputEchoReplies         0
Tcp-statistics:
    ActiveOpens               539
    PassiveOpens              10092
    AttemptFails              36
    EstablishedResets         8505
    CurrentEstab              25
    InputSegments             82339095
    OutputSegments            78020025
    NumRetransSegments        973
    InputErrs                 0
    OutputResets              8592
Udp-statistics:
    InputDatagrams            727256
    InputErrors               0
    OutputDatagrams           727256
    RecvbufErrors             0
    SendbufErrors             0
Extended-IP-statistics:
    InputMcastPkts            3229355
    OutputMcastPkts           278757
    InputBcastPkts            253514
    OutputBcastPkts           0
(END)
```

show system routing

Shows system routing table information.

Syntax:

```
show system routing table
```

Command mode:

All

Example:

SDN-VE @SDN-VE-Controller > show system routing table			
Destination	Gateway	Device	Type
9.121.62.0	0.0.0.0	eth0	UP
default	9.121.62.1	eth0	UP GW

show system uptime

Shows information about the duration the system has been up.

Syntax:

show system uptime

Command mode:

All

Example:

SDN-VE @SDN-VE-Controller > show system uptime	
Time:	12:22:49
Uptime:	21 days, 1 hour, 38 minutes, 16 seconds
Load average:	0.00, 0.00, 0.00

system sdn-ve log

Shows SDN VE log file contents.

Syntax:

system sdn-ve log cat log-file <log file name> [stream]

Command mode:

All

Example:

SDN-VE @SDN-VE-Controller > system sdn-ve log cat log-file log1 stream	
--	--

system sdn-ve log

Lists all SDN VE log files.

Syntax:

```
system sdn-ve log list
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > system sdn-ve log list
```

Date	File size	Logfile name
Jun 18 10:38	0	audit.log
Jun 30 12:34	1.3M	web_access_log_2014-06.txt
Jun 30 23:20	1.9M	opendaylight-2.log.gz
Jul 3 11:40	1.7M	opendaylight_SB-1.log.gz
Jul 9 12:31	39.6K	techdump.0
Jul 13 12:39	1.8M	opendaylight-1.log.gz
Jul 15 06:06	1.7M	web_access_log_2014-07.txt
Jul 15 06:17	39.1M	opendaylight.sb.log
Jul 15 06:17	6.9M	opendaylight.log

system sdn-ve status

Shows SDN VE status.

Syntax:

```
system sdn-ve status
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > system sdn-ve status

SDN-VE is RUNNING
```

Running Configuration Information

show tech

Show system technical information. This command consolidates the following system, configuration, and run time information from various other show commands:

- SDN VE Version
- Cluster Configuration
- High-Availability Configuration
- Host Configuration
- Topology
- Log Levels
- Flows and Flowsets
- LDAP and RADIUS server Configuration
- Tenant Configuration
- Switch Information
- Connectivity Group Configuration
- Subnet Configuration
- Port and Port Group Configuration
- System Configuration
- NIST Configuration
- Replication and Monitoring Configuration
- Layer 3 Configuration
- Policy Configuration
- Waypoint Connectivity Service Configuration

Please note that the output may exceed 10K of data, depending on your configuration.

If you want to capture the data to a file, such as for support or diagnostic purposes, set the communication software on your workstation to capture session data prior to issuing the command.

Syntax:

```
show tech dump support [file]
```

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show tech dump support file  
Tech dump output saved to file: techdump.0
```

Tenant Information

show tenant

Shows information about all tenants or a particular tenant.

Syntax:

show tenant [id <tenant ID> | name <tenant name>]

Command mode:

All

Example:

SDN-VE @SDN-VE-Controller > show tenant				
Id	Name	Domain_Type	Replication factor	Description
1	DOVE ADMIN	DOVE	2	Admin Tenant for DOVE, Created at startup
10	xyz	OF	0	
100	dove_test	DOVE	2	
1000	SDN9990104_T1	OF	0	-
2	OF ADMIN	OF	0	Admin Tenant for OF, Created at startup

Topology Information

show topology

Shows information about linked source and destination nodes and ports.

Syntax:

show topology switch [discoveredLinks | userConfiguredLinks]

Command mode:

Privileged Executive

Example:

SDN-VE @SDN-VE-Controller > show topology switch discoveredLinks				
Edge Details are as shown below..				
Name	Src Node Id	Src Port Id	Dst	
Node Id	Dst Port Id	Latency	Bandwidth	

s1-eth1	00:00:00:00:00:00:01	1		
00:00:00:00:00:00:02	3	-	10000000000	
s3-eth3	00:00:00:00:00:00:03	3		
00:00:00:00:00:00:01	2	-	10000000000	
s2-eth3	00:00:00:00:00:00:02	3		
00:00:00:00:00:00:01	1	-	10000000000	
s1-eth2	00:00:00:00:00:00:01	2		
00:00:00:00:00:00:03	3	-	10000000000	

SDN-VE @SDN-VE-Controller > show topology switch userConfiguredLinks			

Name	Src Node	Connector	Dst Node
Connector	Status		

System Upgrade Information

show upgrade

Shows whether or not a database upgrade is required.

Syntax:

show upgrade status

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show upgrade status

9.126.149.2 : Thu Jul 31 05:01:57 GMT 2014 : Rolling-Upgrade initiated
(current version : 2.0.0.4.2.736)
9.126.149.3 : Thu Jul 31 05:02:04 GMT 2014 : Downloading new image
9.126.149.3 : Thu Jul 31 05:02:12 GMT 2014 : Image verification pass (new
version : 2.0.0.4.2.740)
9.126.149.3 : Thu Jul 31 05:02:12 GMT 2014 : Leaving Cluster
9.126.149.3 : Thu Jul 31 05:02:12 GMT 2014 : Stopping
southbound/northbound communication
9.126.149.3 : Thu Jul 31 05:02:12 GMT 2014 : Executing upgrade script
9.126.149.3 : Thu Jul 31 05:03:43 GMT 2014 : Upgrade complete
9.126.149.3 : Thu Jul 31 05:03:43 GMT 2014 : Upgrading node 9.126.149.2
9.126.149.2 : Thu Jul 31 05:03:51 GMT 2014 : Image verification pass (new
version : 2.0.0.4.2.740)
9.126.149.2 : Thu Jul 31 05:03:52 GMT 2014 : Leaving Cluster
9.126.149.2 : Thu Jul 31 05:03:52 GMT 2014 : Stopping
southbound/northbound communication
9.126.149.2 : Thu Jul 31 05:03:52 GMT 2014 : Executing upgrade script
9.126.149.2 : Thu Jul 31 05:05:24 GMT 2014 : Upgrade complete
```

Users Information

show users

Shows information about users that have access to the SDN VE setup.

Syntax:

show users

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show users

User                Role                Tenant
-----
admin               System-Admin        System-Admin
```

SDN VE Version Information

show version

Shows information about the SDN VE controller version currently installed on the system.

Syntax:

show version

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller > show version

SDN-VE version = 2.0.0
Build version   = 4.2.592
```

Connectivity Group Policy Information

show cgpolicy

Shows information about policies configured for connectivity groups.

Syntax:

show cgpolicy

Command mode:

Global Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config)# show cgpolicy

Group1 Id       : 6cecd2d5-9344-4aba-a82d-3a596c294c05
Group2 Id       : f2b018c5-b481-4e6b-b296-390e8b0aa534
Policy Id       : b9de0f0a-86cf-4605-90fa-f273fbd42543
Traffic type    : UNICAST
Directional     : BI_DIRECTIONAL
```

NIST Information

show nist-mode

Shows information about NIST configuration: enabled or disabled.

Syntax:

show nist-mode

Command mode:

Global Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config)# show nist-mode

NIST-mode status: disable
```

DOVE Configuration Information

Shows SDN VE overlay configuration information.

show sdnve-dove dcs-stats

Shows statistical information about the configured DCS appliances.

Syntax:

show sdnve-dove dcs-stats node-id <Node ID> tenant-id <Tenant ID>

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove dcs-stats node-id 1
tenant-id 1

DCS Statistics
-----
policy_lookup_count      :0
multicast_lookup_count   :0
endpoint_update_rate     :0
internal_gw_lookup_count :2824
policy_lookup_rate       :0
endpoint_lookup_count     :8571
endpoint_update_count     :2724
endpoint_lookup_rate      :0
```

show sdnve-dove dgw-interfaces

Shows information about the gateways configured on the overlay network.

Syntax:

show sdnve-dove dgw-interfaces id <gateway ID>

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove dgw-interfaces id 2
```

show sdnve-dove domain-separation

Shows information about the configured domains.

Syntax:

```
show sdnve-dove domain-separation
```

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove domain-separation  
Separation: no
```

show sdnve-dove external-ip

Shows information about the external IP address of the SDN VE controller.

Syntax:

```
show sdnve-dove external-ip
```

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove external-ip  
external ip 9.121.62.240 mask 255.255.254.0
```

show sdnve-dove gateway-sessions

Shows information about the sessions on a gateway.

Syntax:

```
show sdnve-dove gateway-sessions gw-index <gateway index> type  
{outbound | internal | dynamic}
```

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove gateway-sessions  
gw-index 1 type outbound
```

show sdnve-dove ha

Shows information about the high-availability configuration.

Syntax:

```
show sdnve-dove ha
```

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove ha  
  
ha            running  
external ip 9.121.62.240 mask 255.255.254.0  
primary  ip 9.121.62.118 mask 255.255.254.0  
secondary ip 9.121.62.116 mask 255.255.254.0
```

show sdnve-dove peers

Shows information about the peers configured for high-availability.

Syntax:

```
show sdnve-dove peers
```

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove peers  
  
primary ip 9.121.62.118 mask 255.255.254.0  
secondary ip 9.121.62.116 mask 255.255.254.0
```

show sdnve-dove service-appliances

Show all Distributed Services Appliances (DSAs). Each type of DSA (DCS and DGW) are shown in separate tables. If a DCS has a role assigned, it is marked as Y in the information table. Otherwise as N.

Syntax:

show sdnve-dove service-appliances

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove service-appliances
```

DCS Service Appliances:						
ID	IP	SERVICE CAPABILITY	ROLE ASSIGNED	AGE_TIME	CONFIG VERSION	BUILT VERSION
1	9.121.62.178	CS	N	7 s	0/ 15	1.0.0.131007

GW Service Appliances:						
ID	IP	SERVICE CAPABILITY	ROLE ASSIGNED	AGE_TIME	CONFIG VERSION	BUILT VERSION
1	9.121.62.178	GW	Y	5 s	15/ 15	1.0.0.131007

show sdnve-dove SNAT-pool-size

Shows information about the configured SNAT pool size.

Syntax:

show sdnve-dove SNAT-pool-size

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove SNAT-pool-size
```

SNAT_pool_size: 1

show sdnve-dove switch-info

Shows information about the configured switches.

Syntax:

show sdnve-dove switch-info

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove switch-info

Tunnel Endpoint IP
=====
200.1.1.13
200.1.1.14
```

show sdnve-dove sync-status

Shows information about the HA synchronization status.

Syntax:

```
show sdnve-dove sync-status
```

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove sync-status

Synchronization-status   Synchronized
```

show sdnve-dove syslog

Shows information about the logging configuration.

Syntax:

```
show sdnve-dove syslog
```

Command mode:

SDN VE DOVE Configuration

Example:

SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove syslog			
Module	Log-Level	Log-Console	Log-Flag
-----	-----	-----	-----
dsw	INFO	Disable	Enable
raw	INFO	Disable	Enable
dps	INFO	Disable	Enable
vrnmg	INFO	Disable	Enable
dgw	INFO	Disable	Enable
sys	INFO	Disable	Enable

show sdnve-dove underlay-network

Shows information about the underlay network settings.

Syntax:

```
show sdnve-dove underlay-network
```

Command mode:

SDN VE DOVE Configuration

Example:

SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove underlay-network			
ID	NET	MASK	NEXTHOP
1	9.121.62.0	255.255.255.254.0	9.121.62.1

show sdnve-dove vrrp

Shows information about the VRRP high-availability configured for the gateway.

Syntax:

```
show sdnve-dove vrrp id <VRRP ID>
```

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove vrrp id 2
```

show sdnve-dove vxlan-port

Shows information about the VXLAN port configured.

Syntax:

```
show sdnve-dove vxlan-port
```

Command mode:

SDN VE DOVE Configuration

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# show sdnve-dove vxlan-port  
VXLAN PORT: 8472
```

Chapter 20. Configuration Commands

This chapter discusses how to use the individual CLI for making configuration changes.

Global Configuration Mode

Cluster Configuration

cluster disconnect

Disconnects the cluster.

Syntax:

```
cluster disconnect
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# cluster disconnect
```

cluster rejoin

Rejoins a cluster.

Syntax:

```
cluster rejoin
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# cluster rejoin
```

cluster name

Configure a primary node in the cluster. A primary node acts as a rendezvous point for other nodes in the cluster. This configuration command should be executed on all the nodes in the cluster, including the primary node(s). All nodes must have identical configuration. Multiple nodes can be designated as primary nodes.

Syntax:

```
cluster name <name> node-list ip <IP address(es)>
```

Parameters:

<name> Name of the cluster.

<IP address(es)> IP address of the primary node. If multiple primary nodes are configured, specify the IP addresses separated by a comma.

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# cluster name cluster1 node-list ip
192.168.1.1,192.168.1.4
```

Flowset Configuration

flowset add

Configure a flow group.

Syntax:

```
flowset add name <name> [priority <priority value>]
[ idle <idle timeout value>] [hard <hard timeout value>] [install]
```

Parameters:

<name> Name of the flow set.

<priority value> (Optional). Priority value of the flow set (0-65535 seconds).

<idle timeout value> (Optional). Idle timeout value (0-65535 seconds). If no match is found for the flow set for the configured time, the flow set is removed from the table.

<hard timeout value> (Optional). Hard timeout value (0-65535 seconds). Flow set is removed from the table after the configured time irrespective of the match status.

[install] (Optional). Install the flow set.

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# flowset add name fs1 priority 2 idle 5 hard 10
install
```

no flowset add

Delete a flow group.

Syntax:

```
no flowset add name <name> [install]
```


Parameters:

<name> Name of the flow set.
[install] (Optional). Remove the flow set.

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# flowset add name fs1 priority 2 idle 5 hard 10  
install
```

flowset name

Enter flow group configuration mode.

Syntax:

flowset name *<name or ID>*

Parameters:

<name or ID> Name or ID of the flow set.

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# flowset name fs1  
SDN-VE @SDN-VE-Controller (config-flowset-fs1)#
```

LDAP Server Configuration

Note: To upload an LDAP certificate, use the controller GUI:
<https://<Controller HA External IP address>:8443>

ldap enable

Enable LDAP server.

Syntax:

ldap enable

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# ldap enable
```

ldap server

ldap server domain

Configure LDAP server domain name.

Syntax:

```
ldap server domain name <name>
```

Parameters:

<name> Name for the LDAP server.

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# ldap server domain name mydomain
```

ldap server primary

Configure primary LDAP server.

Syntax:

```
ldap server primary ip <IP address> [port <port number>]
```

Parameters:

<IP address> IP address of primary LDAP server.

<port number> Primary LDAP server port.

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# ldap server primary 9.121.52.45
```

no ldap enable

Disable LDAP server.

Syntax:

```
no ldap enable
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# no ldap enable
```

no ldap server primary

Delete primary LDAP server.

Syntax:

no ldap server primary

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# no ldap server primary
```

Log Setting Configuration

log

Configure logging level for an SDN VE service.

Syntax:

log level *<logging level>* [logger *<logger name>*]

Parameters:

<logging level> Type of log to be generated.

echo

info

trace

debug

warning

error

<logger name> Service for which the log needs to be generated.

loggingBridge

protocol_plugin_openflow

arphandler

broadcast

clustering

commons

dove_config

topology

flowgroupsmanager

flow_reaper

hosttracker

infinispan
interface_manager
layer2
layer3
license
lnp_topology
logical_groups
multicast
nist
odl_services
ofp_processor
ping
pipeline
policymanager
proxy
replication
restore
root
sal
script_interface
security
smarttime
usermanager
web
waypoint

Command mode:
Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# log level info logger topology
```

Multicast Configuration

multicast query-interval

Configure query interval for the Querier.

Syntax:

```
multicast query-interval <30-1800 seconds>
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# multicast query-interval 50
```

multicast vlan

Configure multicast VLAN parameters.

Syntax:

```
multicast vlan <VLAN range>
```

Parameters:

<VLAN range> Range of VLANs (separated by a hyphen) to be used by the forwarding database. Example: 1000-2000.

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# multicast vlan 300-400
```

no multicast vlan

Delete multicast VLAN.

Syntax:

```
no multicast vlan
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# no multicast vlan
```

NIST Configuration

nist

Enable/disable NIST compliance.

Syntax:

```
nist {enable | disable}
```

default setting - disabled

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# nist enable
```

Pagination Configuration

pagination enable/disable

Enable or disable pagination. When enabled, command outputs are displayed in pages.

Syntax:

```
pagination {enable | disable}
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# pagination enable
```

RADIUS Server Configuration

radius enable

Enable RADIUS server.

Syntax:

```
radius enable
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# radius enable
```

radius server primary

Configure primary RADIUS server.

Syntax:

```
radius server primary ip <IP address> key <RADIUS server key>
[port <port number>]
```

Parameters:

<IP address> IP address of primary RADIUS server.
<RADIUS server key> Secret key of the RADIUS server (1-32 characters).
<port number> Primary RADIUS server port.

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# radius server primary ip 192.168.1.1 key
9876345321 port 10330
```

no radius enable

Disable RADIUS server.

Syntax:

```
no radius enable
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# no radius enable
```

no radius server primary

Delete a primary RADIUS server.

Syntax:

```
no radius server primary
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# no radius server primary
```

Reset User Password

reset

Reset user password.

Syntax:

```
reset user name <user name>
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# reset user name usr1
Password:
Confirm Password:
```

SDN VE Configuration

sdnve-dove

Enter SDN VE DOVE configuration mode.

Syntax:

```
sdnve-dove terminal
```

Command mode:

Global Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config)# sdnve-dove terminal
SDN-VE @SDN-VE-Controller (config-sdnve-dove)#
```

See [“SDN VE DOVE Configuration Mode Commands” on page 274](#) for commands available in this mode.

Switch Configuration

switch update

Updates the switch name.

Syntax:

```
switch update dpid <dpid> name <name>
```

Parameters:

<dpid>	Switch DPID
<name>	Switch name

Command mode:
Global Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config)# switch update dpId 00:00:00:00:00:00:01 name  
mySwitch
```

Tenant Configuration

tenant add

Create a tenant.

Syntax:

```
tenant add id <tenant_id> name <tenant_name> type {openflow | dove}  
[descr <description>] [repfactor <n>]
```

Parameters:

<tenant_id>	Tenant ID
<tenant_name>	Tenant name
<description>	Tenant description
<repfactor>	Replication factor.

The replication factor represents the number of DCS nodes on which the system will attempt to copy the tenant configuration. At least two nodes on different hosts are required for HA resilience.

Command mode:
Global Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config)# tenant add id 3456 name tenant21 type dove descr  
IT-Dept repfactor 2  
Tenant created with UUID = 3456
```

tenant id

Enter tenant configuration mode.

Syntax:

```
tenant id <tenant_id>
```

Command mode:
Global Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config)# tenant id 3456
SDN-VE @SDN-VE-Controller(config-tenant-3456)#
```

tenant update

Update a tenant.

Syntax:

```
tenant update id <id> [name <tenant_name>] [repfactor <replication_factor>]
[descr <description>]
```

Parameters:

<ID> Tenant ID.

<tenant_name> Name of the tenant.

<replication_factor>

Replication factor. Number of nodes to which the configuration has to be applied.

The system will attempt to find the requested number of nodes to meet the new replication factor. If the current cluster is not able to meet the new replication factor (for instance, if the replication factor is 4 but only 3 nodes are available), the system will track node availability and perform additional replication when new nodes become available to handle the tenant.

<description> Description of the tenant.

Command mode:

Global Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config)# tenant update id 3456 name tenant2 descr HR-Dept
```

no tenant add

Delete a tenant.

Syntax:

```
no tenant add id <tenant_id>
```

Command mode:

Global Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config)# no tenant add id 3456
```

User Configuration

user add

Creates a user.

Syntax:

```
user add name <name of the user>  
role {tenant-admin | tenant-operator | system-admin}  
tenant id <tenant ID>
```

Note: The tenant-id is required only for roles tenant-admin and tenant-operator.

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller(config)# user add name cmdADMus role tenant-operator  
tenant id 4  
Password:  
Confirm Password:
```

user update

Updates the role of the specified user.

Syntax:

```
user update name <name of the user> role {tenant-admin | tenant-operator |  
system-admin} tenant id <tenant id>
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller(config)# user update name cmdADMus role system-admin  
tenant id 7  
Warning: Any role change will require user to log out and log in  
SDN-VE(config)#
```

no user add

Delete a user.

Syntax:

no user add name *<name of the user>* Command mode:
Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller(config)# no user add name cmdADMus
```

System Configuration Commands

system cli timeout

Any CLI session will be automatically logged out if idle for the length of time set.

Syntax:

```
system cli timeout [<minutes>]
```

Parameters:

<minutes> (Optional). Timeout period in minutes (1-30). Default value is 10 minutes.

Command mode:

Privileged Exec and above.

Example:

```
SDN-VE @SDN-VE-Controller> system cli timeout 8
```

system dateortime

View or configure the System date or time.

Syntax:

```
system dateortime {date|time} [set <date or time>]
```

Parameters:

<date or time> Date - in dd/mm/yyyy format.
Time - in hh:mm format.

Command mode:

Privileged Exec and above.

Example:

```
SDN-VE @SDN-VE-Controller> system dateortime date  
15 Jan 2014
```

```
SDN-VE @SDN-VE-Controller> system dateortime time set 12:51
```

```
System time will require a restart of all nodes in the cluster.  
WARNING: Previous network state will not be preserved.  
Continue (y/n)? n
```

system domain-separation

Enable or disable separation of domains.

Syntax:

```
system domain-separation [true | false]
```

Command mode:

SDN VE DOVE Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# system domain-separation true
```

system hostname

View or configure the system host name.

Syntax:

```
system hostname [<name>]
```

Parameters:

<name> Host name.

Command mode:

Privileged Exec and above.

Example:

```
SDN-VE @SDN-VE-Controller> system hostname  
Hostname: SDN-VE-Controller
```

system ipmgmt

system ipmgmt ip

Set a static controller address by specifying an IPv4 address and mask or CIDR designation.

(or)

Enable DHCP IPv4 address configuration.

Syntax:

```
system ipmgmt ip static <IPv4 address> <netmask>
system ipmgmt ip static <CIDR>
system ipmgmt ip dhcp
```

Parameters:

<IPv4 address>	IPv4 address in dotted decimal notation (a.b.c.d)
<netmask>	IPv4 netmask in dotted decimal notation (a.b.c.d)
<CIDR>	IPv4 address in cidr format (a.b.c.d/e)

Command mode:

Global Configuration mode

Example:

SDN-VE @SDN-VE-Controller(config)# system ipmgmt ip static 10.10.0.10 255.255.255.0
SDN-VE @SDN-VE-Controller(config)# system ipmgmt ip static 10.10.0.10/24
SDN-VE @SDN-VE-Controller(config)# system ipmgmt ip dhcp

system ipmgmt nameserver

Configure a name server by specifying an IPv4 address and mask or CIDR designation.

Syntax:

```
system ipmgmt nameserver <IPv4 address> <netmask>
system ipmgmt nameserver <CIDR>
```

Parameters:

<IPv4 address>	IPv4 address in dotted decimal notation (a.b.c.d)
<netmask>	IPv4 netmask in dotted decimal notation (a.b.c.d)
<CIDR>	IPv4 address in cidr format (a.b.c.d/e)

Command mode:

Global Configuration mode

Example:

SDN-VE @SDN-VE-Controller(config)# system ipmgmt nameserver 10.10.0.10 255.255.255.0
SDN-VE @SDN-VE-Controller(config)# system ipmgmt nameserver 10.10.0.10/24

system ipmgmt nexthop

Set the controller gateway address via specifying an IPv4 address and mask or CIDR designation.

Syntax:

```
system ipmgmt nexthop <IPv4 address><netmask>
system ipmgmt nexthop <CIDR>
```

Parameters:

<IPv4 address>	IPv4 address in dotted decimal notation (a.b.c.d)
<netmask>	IPv4 netmask in dotted decimal notation (a.b.c.d)
<CIDR>	IPv4 address in CIDR format (a.b.c.d/e)

Command mode:

Global Configuration mode

Example:

SDN-VE @SDN-VE-Controller(config)# system ipmgmt nexthop 10.10.0.1 255.255.255.0
SDN-VE @SDN-VE-Controller(config)# system ipmgmt nexthop 10.10.0.1/24

no ipmgmt nameserver

Delete a name server.

Syntax:

```
no ipmgmt nameserver id <ID>
```

Parameters:

<ID>	Nameserver ID.
------	----------------

Command mode:

Global Configuration mode

Example:

SDN-VE @SDN-VE-Controller(config)# no ipmgmt nameserver id 1
--

no ipmgmt nexthop

Delete a next hop.

Syntax:

```
no ipmgmt nexthop
```


Command mode:
Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller(config)# no ipmgmt nexthop
```

system license

Add License.

Syntax:
`system license <license_key>`

Parameters:
`<license_key>` 64 character license key

Command mode:
Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller(config)# system license

Please enter your license (64 chars): <license_key>

Pass through the complete license agreement and type in ":q" to quit the
license agreement.

In the prompt "Enter 'y' to accept or 'n' to disagree":- Input 'y' to
accept and add the license
```

Note: The license could be either added afresh or upgraded from Temporary to Permanent validity, whereas cannot be deleted or downgraded.

system reboot

Restart the controller VM. Controller operation is temporarily halted while the software is restarted. When the reboot is complete, the saved configuration is restored and normal operation is resumed.

Note: If HA has been configured between the primary and secondary controllers, you must run the command `system reboot` on both the primary and secondary controllers. Reboot the secondary controller only after the primary controller GUI comes up after the reboot.

Syntax:
`system reboot`

Command mode:
Privileged Exec and above.

Example:

```
SDN-VE @SDN-VE-Controller> system reboot
```

system restart

Restart the affected processes when configuration is changed. When the restart is complete, the saved configuration is restored and normal operation is resumed. This command is required when changing the enable or disable status of NIST and PKI configuration.

Note: If HA has been configured between the primary and secondary controllers, you must run the command `system restart` on both the primary and secondary controllers. Reboot the secondary controller only after the primary controller GUI comes up after the restart.

Syntax:

```
system restart
```

Command mode:

Privileged Exec and above.

Example:

```
SDN-VE @SDN-VE-Controller> system restart
```

system sdn-ve

SDN VE system configuration.

system sdn-ve log rm

Delete a log file.

Syntax:

```
system sdn-ve log rm file <log file name>]
```

Parameters:

<log file name> Name of the log file.

Command mode:

Privileged Exec and above.

Example:

```
SDN-VE @SDN-VE-Controller> system sdn-ve log rm file audit.log
```

system sdn-ve authenticate

Authenticate with SDN VE.

Syntax:

```
system sdn-ve authenticate
```

Command mode:

Privileged Exec and above.

Example:

```
SDN-VE @SDN-VE-Controller> system sdn-ve authenticate
```

system shutdown

Shutdown the controller.

Syntax:

```
system shutdown
```

Command mode:

Privileged Exec and above.

Example:

```
SDN-VE @SDN-VE-Controller> system shutdown
```

system SNAT-pool-size

Specify the number of IP addresses for static network address translation (NAT).

Syntax:

```
system SNAT-pool-size <number of IP addresses>
```

Command mode:

SDN VE DOVE Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# system SNAT-pool-size 8
```

system syslog

Configure syslog options for modules running on the controller.

Log Level

Set the log level for a specific module (process).

Logging messages are stored in the /flash/dmc_syslog.log file.

Syntax:

```
system syslog <module name><log level> [enable]
```

Parameters:

<module name> The name of the target module:

- dps
- sys
- dgw
- dsw
- vrmgr
- raw

<log level> The log level of the target module:

- EMERGENCY
- ALERT
- CRITICAL
- ERROR
- WARNING
- NOTICE
- INFO
- DEBUG

Command mode:

SDN VE DOVE Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# system syslog dpsdebug ALERT
```

Syslog Enable or Disable

System Log Module Control

You can enable or disable logging for specific modules (processes). Logging is disabled by default.

Syntax:

```
system syslog <module>{enable|disable}
```

Parameters:

<module>

The name of the target module:

- dps
- sys
- dgw
- dsw
- vrMgr
- raw

Command mode:

SDN VE DOVE Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# system syslog dps enable
```

Console Log

Enable logging to console.

Syntax:

```
system syslog <module> console {enable|disable}
```

Parameters:

<module>

The name of the target module:

- dps
- sys
- dgw
- dsw
- vrMgr
- raw

Command mode:

SDN VE DOVE Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# system syslog dps console enable
```

system upgrade

Upgrade the controller software version.

Note: The upgrade image should be placed on a web server that is accessible to the controller. The image file extension will be `.img`.

Syntax:

```
system upgrade <URL>
```

Command mode:

Privileged Executive and above.

Parameters:

<URL> Location of the image file.

Example:

```
SDN-VE @SDN-VE-Controller> system upgrade  
http://9.121.62.106/TestSetup/ibm-sdn-ve-upgrade-4.2.740.img
```

system vxlan

Set the VXLAN UDP port number.

Syntax:

```
system vxlan port {4789|8472}
```

Command mode:

SDN VE DOVE Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# system vxlan port 4789
```

Flow Group Configuration Mode Commands

The commands in this section apply to flow group configuration mode. See [“flowset name” on page 233](#) for steps to access this mode.

flow add
flow update

Add a flow to the flow set.

Update a flow that is part of a flow set.

Syntax:

```
flow {add | update} name <name> switch <switch ID>
priority <priority value>]
{idle <idle timeout value>}
{hard <hard timeout value>}
{action <action>}
[match <match action>]
```

Parameters:

<name>	Name of the flow.
<switch ID>	Format: xx:xx:xx:xx:xx:xx:xx:xx
<priority value>	(Optional). Priority value of the flow set (0-65535 seconds).
<idle timeout value>	(Optional). Idle timeout value (0-65535 seconds). If no match is found for the flow for the configured time, the flow is removed from the table.
<hard timeout value>	(Optional). Hard timeout value (0-65535 seconds). Flow is removed from the table after the configured time irrespective of the match status.
<action>	Select action for the flow. You may specify multiple actions from the following list:
drop	Drop the flow.
strip-vlan	Remove VLAN tag.
output <out port(s)>	Out port(s) for the flow.
set-vlan-id <VLAN ID (0-4094)>	Out VLAN ID for the flow. 0 to be used for untagged packets.
set-dl-src <MAC address>	Source MAC address for the flow.
set-dl-dst <MAC address>	Destination MAC address for the flow

`set-dl-type <string>`
 Destination type for the flow

`set-vlan-pcp <0-7>`
 VLAN priority for the flow.

`set-nw-tos <0-63>`
 Type of service for the flow. 0 - lowest priority; 63 - highest priority.

`set-nw-src <IP address>`
 Network source for the flow.

`set-nw-dst <IP address>`
 Network destination for the flow.

`set-tp-src <integer>`
 Transport layer source for the flow.

`set-tp-dst <integer>`
 Transport layer destination for the flow.

`<match action>` Select match criteria for the flow. You may specify multiple match criteria from the following list:

`in-port <port number>`
 Ingress port for the flow.

`dl-type <string>`
 Ethernet type for the flow.

`dl-src <MAC address>`
 Source MAC address for the flow.

`dl-dst <MAC address>`
 Destination MAC address for the flow.

`dl-vlan <VLAN ID (1-4095 or 65535)>`
 Destination VLAN ID for the flow; 65535 for untagged packets.

`dl-vlan-pcp <0-7>`
 VLAN priority for the flow.

`nw-src <IP address>`
 Source IP of the flow.

`nw-dst <IP address>`
 Destination IP address for the flow.

`nw-tos-dscp <0-63>`
 Type-of-Service for the flow.

`nw-protocol <0-255>`
 IP protocol for the flow.

`tcp-src <integer>`
 Transport layer source for the flow.

tcp-dst <integer>

Transport layer destination for the flow.

Command mode:

Flow Group Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-flowset-fs1)# flow add name a1 switch  
00:00:00:01:03:04:05:06 priority 2 action drop match nw-tos-dscp 5
```

no flow add

Delete a flow from the flow set.

Syntax:

no flow add name <name>

Command mode:

Flow Group Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-flowset-fs1)# no flow add name a1
```

Tenant Configuration Mode Commands

The commands in this section apply to tenant configuration mode. See [“tenant id” on page 241](#) for steps to access this mode.

Connectivity Group Configuration

group add

Create a connectivity group.

Syntax:

```
group add name <CG name> admin-state <CG status> [vnid <ID>]
[traffic <traffic type>] [precedence <level>] [limitDelay <value>]
[limitThroughput <value>] [limitReliability <value>]
[average_rate <value in KBps>] [peak_rate <value in KBps>]
[burst_rate <value in KiloBytes>] [id <group ID>]
[group-type {dedicated | shared | external}] [isNeutron {true | false}]
```

Parameters:

<CG name>	Name of the connectivity group.
<CG status>	Administrative state of the connectivity group: up or down. If down, the group does not forward packets.
vnid <ID>	UUID of the group (1-36 alphanumeric characters).
<traffic type>	Traffic type of the group: BEST_EFFORT BACKGROUND EXCELLENT_EFFORT CRITICAL_APPLICATIONS VIDEO VOICE INTERNETWORK_CONTROL NETWORK_CONTROL
<level>	Precedence level: ROUTINE PRIORITY IMMEDIATE FLASH FLASH_OVERRIDE CRITIC_ECP INTERNETWORK_CONTROL NETWORK_CONTROL
limitDelay <value>	integer value: 0 or 1 0 - False 1- True

limitThroughput <value>
 integer value: 0 or 1
 0 - False
 1- True

limitReliability <value>
 integer value: 0 or 1
 0 - False
 1- True

average_rate <value in KBps>
 The average number of kilobytes per second (KBps) to allow across a port or a portgroup.

peak_rate <value in KBps>
 The number of kilobytes per second (KBps) to allow across a port or a portgroup, when it is sending/receiving a burst of traffic.

burst_rate <value in KiloBytes>
 Maximum number of kilobytes to allow in a burst.

<group ID> UUID of the group (1-36 alphanumeric characters).

group-type Group resource type: dedicated/shared/external.

isNeutron Does the Group use OpenStack Neutron APIs? Values: true or false.

Command mode:
 Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# group add name test admin-state up
group-type dedicated
```

group update

Update connectivity group specifications.

Syntax:

```
group update id <group ID> [name <CG name>] [admin-state <CG status>]
[traffic <traffic type>] [precedence <level>] [limitDelay <value >]
[limitThroughput <value >] [limitReliability <value >]
[average_rate <value in KBps >] [peak_rate <value in KBps >]
[burst_rate <value in KiloBytes >]
```

Parameters:

<group ID> UUID of the group (1-36 alphanumeric characters).
 <CG name> Name of the connectivity group.

<CG status> Administrative state of the connectivity group: up or down.
If down, the group does not forward packets.

<traffic type> Traffic type of the group:
BEST_EFFORT
BACKGROUND
EXCELLENT_EFFORT
CRITICAL_APPLICATIONS
VIDEO
VOICE
INTERNETWORK_CONTROL
NETWORK_CONTROL

<level> Precedence level:
ROUTINE
PRIORITY
IMMEDIATE
FLASH
FLASH_OVERRIDE
CRITIC_ECP
INTERNETWORK_CONTROL
NETWORK_CONTROL

limitDelay *<value>*
integer value: 0 or 1
0 - False
1- True

limitThroughput *<value>*
integer value: 0 or 1
0 - False
1- True

limitReliability *<value>*
integer value: 0 or 1
0 - False
1- True

average_rate *<value in KBps>*
The average number of kilobytes per second (KBps) to allow across a port or a portgroup.

peak_rate *<value in KBps>*
The number of kilobytes per second (KBps) to allow across a port or a portgroup, when it is sending/receiving a burst of traffic.

burst_rate *<value in KiloBytes>*
Maximum number of kilobytes to allow in a burst.

Command mode:
Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# group add name test admin-state up
group-type dedicated
```

group id

Enter group configuration mode.

Syntax:

```
group id <UUID of the group>
```

Command mode:

Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# group id e263696b-3cbf-4ca8-b884-c7
9b9401890b
SDN-VE @SDN-VE-Controller(config-tenant-1-group)#
```

no group add

Delete a connectivity group.

Syntax:

```
no group add id <UUID of the group>
```

Command mode:

Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# no group add id
e263696b-3cbf-4ca8-b884-c79b9401890b
```

Connectivity Group Policy Configuration

cgpolicy add

Configure a policy to be applied between two connectivity groups.

Syntax:

```
cgpolicy add id <id1> id <id2> traffic-type <traffic type> directional
<traffic direction> [snat start-ip <Start IP address> end-ip <End IP address>
start-port <Start port number> end-port <End Port number>]
```

Parameters:

<id1> ID of the first Connectivity Group.

<id2> ID of the second Connectivity Group.

<i><traffic type></i>	Type of traffic: Unicast or Multicast.
<i><traffic direction></i>	Direction of traffic flow between the two groups: UNI_DIRECTIONAL or BI_DIRECTIONAL.
SNAT	If the second connectivity group is of type external, you may specify the IP address and port range to be used for NAT. This configuration is optional.
<i><Start IP address></i>	Starting IP address of the range of addresses you want to allocate for NAT.
<i><End IP address></i>	Ending IP address of the range of addresses you want to allocate for NAT.
<i><Start Port number></i>	Starting port number of the range of ports to be assigned for NAT. Port numbers can be in the range: 1-65535.
<i><End Port number></i>	Ending port number of the range of ports to be assigned for NAT. Port numbers can be in the range: 1-65535.

Command mode:
Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# cgpolicy add id 1 id 2 traffic-type
Unicast directional BI_DIRECTIONAL
```

cgpolicy update

Update a policy applied between two connectivity groups.

Syntax:

```
cgpolicy update id <policy id> dnat ip <IPv4 address> port <DNAT port>
```

Parameters:

<i><policy id></i>	ID of the existing policy.
<i><IPv4 address></i>	DNAT IPv4 address.
<i><port></i>	DNAT port.

Command mode:
Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# cgpolicy update id 1 dnat ip 10.10.10.1
port 80
```

no cgpolicy

Delete a policy applied between two connectivity groups.

Syntax:

```
no cgpolicy id <Policy ID>
```

Command mode:

Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# no cgpolicy id 1
```

Flow Replication Configuration

monitor add

Add a flow replication session.

Syntax:

```
monitor add name <Session name>
{src-ip <IP address> | src-mac <MAC address> | src-any}
{dest-ip <IP address> | dest-mac <MAC address> | dest-any}
{target-ip <IP address> | target-mac <MAC address>} [protocol <type>]
[src-port <Port number or any>] [dest-port <Port number>]
[target-tenant <Tenant name>]
```

Note: You cannot specify both `src-any` and `dest-any`.

Parameters:

<Session name>	Name of the session.
src-ip <IP address>	Source IP address of the flow.
src-mac <MAC address>	Source MAC address of the flow.
src-any	Flow can be from any source.
dest-ip <IP address>	Destination IP address of the flow.
dest-mac <MAC address>	Destination MAC address of the flow.
dest-any	Flow can be to any destination.
target-ip <IP address>	IP address of the replication host.
target-mac <MAC address>	MAC address of the replication host.
protocol <type>	Protocol type: tcp, udp, icmp, any

src-port <Port number or any>
Specify source port number. You can also specify any.

dest-port <Port number or any>
Specify destination port number. You can also specify any.

target-tenant <Tenant name>
Replication tenant name.

Command mode:
Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# monitor name tcp-5 src-ip 10.1.1.1 tcp  
dest-any dest-port 22 target-mac 2a:2b:2c:01:2c:04
```

monitor start

Start a flow replication session.

Syntax:

monitor start name <Session name>

Parameters:

<Session name> Name of the session.

Command mode:
Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# monitor start name S1
```

monitor stop

Stop a flow replication session.

Syntax:

monitor stop name <Session name>

Parameters:

<Session name> Name of the session.

Command mode:
Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# monitor stop name S1
```

monitor update

Update a flow replication session.

Syntax:

```
monitor update name <Session name>
{target-ip <IP address> | target-mac <MAC address>}
[target-tenant <Tenant name>]
```

Parameters:

<Session name> Name of the session.

target-ip <IP address>
 IP address of the replication host.

target-mac <MAC address>
 MAC address of the replication host.

target-tenant <Tenant name>
 Replication tenant name.

Command mode:

Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# monitor update name tcp-5 target-mac
2a:2b:2c:01:2c:04
```

no monitor add

Delete a flow replication session. You can also delete all replication sessions.

Syntax:

```
no monitor add {all | name <Session name> }
```

Parameters:

all Delete all sessions.

<Session name> Name of the session.

Command mode:

Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# no monitor add all
```

Flow Redirection Configuration

redirect add

Add a flow redirection session.

Syntax:

```
redirect add name <Session name>
{src-ip <IP address> | src-mac <MAC address> | src-any}
{dest-ip <IP address> | dest-mac <MAC address> | dest-any}
{target-ip <IP address> | target-mac <MAC address>} [protocol <type>]
[src-port <Port number or any>] [dest-port <Port number>]
[target-tenant <Tenant name>]
```

Note: You cannot specify both `src-any` and `dest-any`.

Parameters:

<code><Session name></code>	Name of the session.
<code>src-ip <IP address></code>	Source IP address of the flow.
<code>src-mac <MAC address></code>	Source MAC address of the flow.
<code>src-any</code>	Flow can be from any source.
<code>dest-ip <IP address></code>	Destination IP address of the flow.
<code>dest-mac <MAC address></code>	Destination MAC address of the flow.
<code>dest-any</code>	Flow can be to any destination.
<code>target-ip <IP address></code>	IP address of the redirection host.
<code>target-mac <MAC address></code>	MAC address of the redirection host.
<code>protocol <type></code>	Protocol type: tcp, udp, icmp, any
<code>src-port <Port number or any></code>	Specify source port number. You can also specify any.
<code>dest-port <Port number or any></code>	Specify destination port number. You can also specify any.
<code>target-tenant <Tenant name></code>	Redirection tenant name.

Command mode:

Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# redirect name tcp-5 src-ip 10.1.1.1 tcp  
dest-any dest-port 22 target-mac 2a:2b:2c:01:2c:04
```

redirect start

Start a flow redirection session.

Syntax:

```
redirect start name <Session name>
```

Parameters:

<Session name> Name of the session.

Command mode:

Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# redirect start name S1
```

redirect stop

Stop a flow redirection session.

Syntax:

```
redirect stop name <Session name>
```

Parameters:

<Session name> Name of the session.

Command mode:

Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# redirect stop name S1
```

redirect update

Update a flow redirection session.

Syntax:

```
redirect update name <Session name>
{target-ip <IP address> | target-mac <MAC address>}
[target-tenant <Tenant name>]
```

Parameters:

<Session name> Name of the session.

target-ip <IP address>
 IP address of the redirection host.

target-mac <MAC address>
 MAC address of the redirection host.

target-tenant <Tenant name>
 Redirection tenant name.

Command mode:

Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# redirect update name tcp-5 target-mac
2a:2b:2c:01:2c:04
```

no redirect add

Delete a flow redirection session.

Syntax:

```
no redirect add name <Session name>
```

Parameters:

<Session name> Name of the session.

Command mode:

Tenant Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1)# no redirect add name tcp-5
```

Subnet Configuration

subnet add

Create a new subnet. This will create a new entry in the controller database that can be later associated with virtual networks (VNIDs).

Syntax:

```
subnet add name <Subnet name> cidr <CIDR IPv4 address> [id <Subnet ID>]
[isNeutron <true or false>] [subnet_type {dedicated | shared | external}]
[gateway <gateway IP address>]
[allocation_pools start <Starting IP address> end <Ending IP address>]
[vlan <VLAN ID>]
```

Parameters:

<Subnet name>	Name of the subnet.
<CIDR IPv4 address>	IPv4 address of the subnet in CIDR format: A.B.C.D/netmask.
<Subnet ID>	UUID of the subnet.
isNeutron	Whether subnet uses OpenStack Neutron APIs or not: true or false.
subnet_type	Three types of subnets are permitted: dedicated or shared or external. Each specific subnet can be associated with only one type at a time.
<gateway IP address>	IP address of the gateway.
<Starting IP address>	Start IP address of the pool of IP addresses allocated for the subnet.
<Ending IP address>	End IP address of the pool of IP addresses allocated for the subnet.
<VLAN ID>	VLAN ID of the subnet. Range: 0-4094. 0 is used for untagged packets.

Command mode:

Tenant Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-tenant-2)# subnet add name Sub2 cidr 2.2.2.2/24 id
2
Subnet created with UUID = 2
```

subnet update

Update specifications of an existing subnet.

Syntax:

```
subnet update id <Subnet ID> [name <Subnet name>]  
[gateway <gateway IP address>] [vlan <VLAN ID>]
```

Parameters:

<Subnet ID>	UUID of the subnet.
<Subnet name>	Name of the subnet.
<gateway IP address>	IP address of the gateway.
<VLAN ID>	VLAN ID of the subnet. Range: 0-4094. 0 is used for untagged packets.

Command mode:

Tenant Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-tenant-2)# subnet update id 2 vlan 100
```

no subnet add

Delete a subnet.

In order to fully remove a subnet from the SDN VE configuration, all DCS nodes (with DCS role shown as Y in the `show sdnve-dove service-appliance` output) must be currently available on the network. If assigned DCS nodes are unavailable, the controller will retain the target subnet information. The configuration elements pertaining to the deleted subnet will not be cleared from the controller until the DCS nodes are made available on the network, or until their DCS roles have been reset using the `no service dcs` command.

Syntax:

```
no subnet add id <Subnet ID>
```

Parameters:

<Subnet ID>	UUID of the subnet.
-------------	---------------------

Command mode:

Tenant Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-tenant-2)# no subnet add id 2
```

Group Configuration Mode Commands

The commands in this section apply to group configuration mode. See [“group id” on page 261](#) for steps to access this mode.

export

Export the specified group configuration to a remote entity.

If the remote entity is a DS 5000V, a profile named `domain.network.vds-name` will be created on the vSwitch, and saved with a VNID set to the specified group ID.

By default, exported profiles will be created on the DS 5000V with 10 ports. Use the CLI present on the DS 5000V to add or remove ports (`config-dvprof mode addports` or `delports`).

Syntax:

```
export ip <vSwitch IP address>
```

Command mode:

Group Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1-group)# export ip 9.121.62.27
```

port add

Note: For DOVE tenants, you can create only a Layer 3 interface.

Add a port to the connectivity group.

Syntax:

```
port add name <port name> admin-state {up | down} mac <MAC address>  
[id <port ID>] ipv4 <IP address>
```

Parameters:

<code><port name></code>	Name of the port.
<code>admin-state</code>	Administrative state of the port: up or down. If down, the port does not forward packets.
<code><MAC address></code>	MAC address of the port. Format: xx:xx:xx:xx:xx:xx
<code><port ID></code>	Port ID.
<code><IP address></code>	IP address for a Layer 3 interface. This associates the port with a subnet.

Command mode:

Group Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1-group)# port add name P1 admin-state up  
mac aa:bb:cc:dd:ee:ff ipv4 1.1.1.3
```

port update

Note: For DOVE tenants, you can update only a Layer 3 interface.

Update port specifications.

Syntax:

```
port update id <port ID> [name <port name>] [admin-state {up | down}]  
[ipv4 <IPv4 address>]
```

Parameters:

<Port ID>	Port ID.
<port name>	Name of the port.
admin-state	Administrative state of the port: up or down. If down, the port does not forward packets.
<IPv4 address>	IPv4 address of the port. The port associates with the subnets using this IP address.

Command mode:

Group Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1-group)# port update id 1 admin-state down
```


subnet

Associate a subnet with the group.

This command will associate a configured subnet with the current group. The command will send a REST message to the DCS modules and to all VLAN gateways and external gateways that are part of the current group.

Syntax:

```
subnet attach id <Subnet ID>
```

Parameters:

<Subnet ID> UUID of the subnet.

Command mode:

Group Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1-group)# subnet attach id 1
```

vlan-gateway

Configure a VLAN gateway for the group. Associate a VLAN to the overlay network (VNID). This enables VLAN to VNID mapping in the selected gateway appliance.

Syntax:

```
vlan-gateway add dgw_id <DGW ID> vlan <VLAN ID>
```

Parameters:

<DGW ID> Index of the Distributed Service appliance (DSA) configured with role of Distributed Gateway (DGW).

<VLAN ID> VLAN ID. Range: 2-4095.

Command mode:

Group Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller(config-tenant-1-group)# vlan-gateway add dgw_id 1 vlan 100
```

SDN VE DOVE Configuration Mode Commands

The commands in this section apply to SDN VE DOVE configuration mode. See [“sdnve-dove” on page 240](#) for steps to access this mode.

external-ip

Set the controller cluster's high-availability (HA) external address by specifying an IPv4 address and mask.

You must use this IP address to access the controller GUI, and not the primary or secondary controller IP addresses.

`https://<Controller HA external IPv4 address>:8443`

Syntax:

`external-ip ip <IP address> mask <netmask>`

Parameters:

`<IP address>` External IP address of the controller.

`<mask>` Subnet mask.

Command mode:

SDN VE DOVE Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# external-ip ip 9.121.62.240 mask 255.255.255.0
```

no external-ip

Delete controller cluster's high-availability (HA) external address.

Syntax:

`no external-ip`

Command mode:

SDN VE DOVE Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# no external-ip
```

ha start

Start the high-availability (HA) feature.

This function will work in the background. When complete, the status will be displayed on the console.

Syntax:

```
ha start
```

Command mode:

SDN VE DOVE Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)#4 ha start
```

ha stop

Stop the high-availability (HA) feature.

This function will work in the background. When complete, the status will be displayed on the console.

Syntax:

```
ha stop
```

Command mode:

SDN VE DOVE Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# ha stop
```

ha-synchronization start

Starts a one-time synchronization of the database from the HA primary node to the HA secondary node.

This function will work in the background. Use the `show ha-synchronization` command to check the status of the synchronization process.

Syntax:

```
ha-synchronization start
```

Command mode:

SDN VE DOVE Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# ha-synchronization start
```

peers

Set the controller HA primary and secondary peer addresses by via specifying an IPv4 address and mask.

Once the peer addresses are set, the internal database will automatically restart before processing can continue.

Syntax:

```
peers primary <IPv4 address> mask <netmask> secondary <IPv4 address>
mask <netmask>
```

Parameters:

<IPv4 address> IPv4 address in dotted decimal notation (a.b.c.d)
<netmask> IPv4 netmask in dotted decimal notation (a.b.c.d)

Command mode:

SDN VE DOVE Configuration mode.

Examples:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# peers primary 9.121.62.116 mask
255.255.255.0 secondary 9.121.62.118 mask 255.255.255.0
```

no peers

Delete controller HA primary and secondary peer addresses.

Syntax:

```
no peers
```

Command mode:

SDN VE DOVE Configuration mode.

Examples:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# no peers
```

service dgw

DOVE Tunnel Endpoint for gateway appliance: This IP address is used to communicate with DOVE Switches and Other DOVE Gateways. This IP address will be used in DOVE Encapsulation headers.

External network IP address for external gateway operation: This IP address is used to communicate with external network.

Syntax:

```
service dgw id <Service appliance ID> add-interface ip <IPv4 address>
mask <netmask> nexthop <gateway IPv4> {dovetunnel|external}
vlan <VLAN ID>
```

Parameters:

<Service appliance ID>	ID of the DGW appliance. Use <code>show sdnve-dove service-appliances</code> command to view the ID.
<IPv4 address>	IPv4 address in dotted decimal notation (a.b.c.d)
<netmask>	IPv4 netmask in dotted decimal notation (a.b.c.d)
<gateway IPv4>	The gateway (next hop) for this interface.
dovetunnel	Add an SDN VE tunnel endpoint IPv4 address. The interface defines an IPv4 address to communicate with SDN VE switches and other SDN VE gateways. This IPv4 address will be used in SDN VE encapsulation headers.
external	Add an external network IPv4 address for external gateway operation. The interface is used to communicate with external networks without SDN VE encapsulation headers.
<VLAN ID>	Specify the VLAN ID (0-4094) of this interface. 0 is used for untagged packets.

Command mode:

SDN VE DOVE Configuration mode.

Examples:

SDN-VE @SDN-VE-Controller (config-sdnve-dove)# service dgw id 1 add-interface ip 1.1.1.10 mask 255.255.255.0 nexthop 1.1.1.254 dovetunnel vlan 0
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# service dgw id 1 add-interface ip 1.1.1.10 mask 255.255.255.0 nexthop 1.1.1.254 dovetunnel vlan 100
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# service dgw id 1 add-interface ip 1.1.1.10 mask 255.255.255.0 nexthop 1.1.1.254 external vlan 0
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# service dgw id 1 add-interface ip 1.1.1.10 mask 255.255.255.0 nexthop 1.1.1.254 external vlan 200

service role dcs

Assign DCS Role to DSA.

Assign a list of Distributed Services Appliances (DSAs) to act as Distributed Connectivity Service (DCS) nodes.

Each DSA can be assigned either a DCS or DGW role. These roles are mutually exclusive. The DCS role can be applied only to DSAs that have no current role assigned. If a target DSA is presently operating in a DGW role, the role must be reset prior to assigning the DCS role (see the `no service role dgw` command).

Syntax:

```
service role dcs ids <DSA list>
```

Parameters:

<i><DSA list></i>	A comma-separated list of target DSA node IDs. IDs are as shown in the <code>show sdnve-dove service-appliance</code> command.
-------------------------	--

Command mode:

SDN VE DOVE Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# service role dcs ids 1,2
```

service role dgw

Assign Gateway Role to DSA.

Assign a registered Distributed Services Appliance (DSAs) to act as a Dove Gateway (DGW) node. Setting this role will allow gateway related configuration on the DSA.

Each DSA can be assigned either a DGW or DCS role. These roles are mutually exclusive. The DGW role can be applied only to DSAs that have no current role assigned. If a target DSA is presently operating in a DCS role, the role must be reset prior to assigning the DGW role (see the `no service role dcs` command).

Syntax:

```
service role dgw ids <DGW list>
```

Parameters:

<i><DGW ID></i>	DGW ID of target gateway appliance. DSA IDs are shown in the <code>show sdnve-dove service-appliance</code> command.
-----------------------	--

Command mode:

SDN VE DOVE Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# service role dgw ids 3,4
```

service gateway id

Enters gateway configuration mode.

Syntax:

```
service gateway id <DGW ID>
```

Parameters:

<DGW ID> DGW ID of target gateway appliance. DSA IDs are shown in the `show sdnve-dove service-appliance` command.

Command mode:

SDN VE DOVE Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# service gateway id 3
SDN-VE @SDN-VE-Controller (config-sdnve-dove-gateway)#
```

no service dcs

Delete a DCS service appliance.

This command can be applied only to a DCS that has no current role assigned: either its underlying DSA has not yet been assigned, or the module has been reset to its basic DSA function, removing the DCS role. To reset the role of a currently assigned DCS prior to deletion, use the `no service role dcs` command.

Syntax:

```
no service dcs id <DCS ID>
```

Parameters:

<DCS ID> DCS ID of target gateway appliance. DSA IDs are shown in the `show sdnve-dove service-appliance` command.

Command mode:

SDN VE DOVE Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# no service gateway dcs id 1
```

no service dgw

Delete a DGW service appliance or delete the IP address of the DGW appliance.

This command can be applied only to a DGW that has no current role assigned: either its underlying DSA has not yet been assigned, or the module has been reset to its basic DSA function, removing the DGW role. To reset the role of a currently assigned DGW prior to deletion, use the `no service role dgw` command.

Syntax:

```
no service dgw id <DGW ID> [ip <IP address>]
```

Parameters:

<DGW ID> DGW ID of target gateway appliance. DSA IDs are shown in the `show sdnve-dove service-appliance` command.

<IP address> IP address of the gateway appliance.

Command mode:

SDN VE DOVE Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# no service dgw id 3
```

no service role

Reset the role of a service appliance.

Syntax:

```
no service role {dgw | dcs} id <DGW or DCS ID>
```

Parameters:

<DGW or DCS ID> DGW or DCS ID of the target gateway appliance. DSA IDs are shown in the `show sdnve-dove service-appliance` command.

Command mode:

SDN VE DOVE Configuration mode.

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# no service role dgw id 3
```


underlay-network subnet

Add a subnet to an underlay network.

Notifies the controller regarding a physical (underlay) IPv4 subnet to which TEPs connect. Nexthop indicates the gateway to be used for reaching TEPs outside the network.

The virtual switches get this configuration when they first connect to the controller and poll the controller every five minutes for updates.

Only one gateway (next hop) can be configured per address/mask pair. To change a gateway, delete the corresponding configuration and create a new one.

This configuration is needed only if TEPs need to communicate with entities outside their network.

Note: The configuration needs to be made before the virtual switches are configured. (Before the TEP VMKNIC is added to the VDS)

If the configuration needs to be changed after vSwitches have been connected, reset the configuration as follows:

1. Make the appropriate changes on the controller and ensure they are correct.
2. Remove the TEP VMKNIC on the Dove Switches and reconnect them. This will trigger another relay of information to and from the DMC and this will update the gateway information on the Dove Switches.

Syntax:

```
underlay-network subnet <IPv4 address> mask <netmask> nexthop  
<gateway IPv4>
```

Parameters:

<IPv4 address>	IPv4 address in dotted decimal notation (a.b.c.d)
<netmask>	IPv4 netmask in dotted decimal notation (a.b.c.d)
<gateway IPv4>	Gateway IPv4 address in dotted decimal notation (a.b.c.d)

Command mode:

SDN VE DOVE Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# underlay-network subnet 10.10.10.0  
mask 255.255.255.0 nexthop 10.10.10.254
```

no underlay-network subnet

Delete underlay network subnet.

Syntax:

```
no underlay-network id <Subnet ID>
```

Command mode:

SDN VE DOVE Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# no underlay-network id 1
```

vrrp add

Add a VRRP router.

Syntax:

```
vrrp add virtual_tep <Virtual TEP IPv4 address>  
virtual_ext_ip <Virtual External IP> gateway index <First gateway index>  
priority <First gateway priority> gateway index <Second gateway index>  
priority <Second gateway priority> virtual_router_id <Virtual router ID>
```

Parameters:

<Virtual TEP IPv4 address>

Virtual IPv4 address in dotted decimal notation (a.b.c.d)

<Virtual External IPv4 address>

Virtual external IPv4 address in dotted decimal notation
(a.b.c.d)

<First gateway index>

Index of the DGW appliance that you want to configure as
master.

<First gateway priority>

Priority of the master DGW appliance. Range: 1-254.

<Second gateway index>

Index of the DGW appliance that you want to configure as
backup.

<Second gateway priority>

Priority of the backup DGW appliance. Range: 1-254.

<First gateway index>

Index of the DGW appliance that you want to configure as
master.

<Virtual router ID> Specify a Virtual Router ID. Range: 1-255. Both DGWs are
configured with this VRID.

Command mode:

SDN VE DOVE Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# vrrp add virtual_tep 1.1.1.10  
virtual_ext_ip 1.1.1.254 gateway index 1 priority 20 gateway index 2 priority 1  
virtual_router_id 1
```

no vrrp

Delete VRRP router.

Syntax:

```
no vrrp id <VRRP HA ID>
```

Command mode:

SDN VE DOVE Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# no vrrp id 1
```

Service Gateway Configuration Mode Commands

The commands in this section apply to service gateway configuration mode. See [“service gateway id” on page 279](#) for steps to access this mode.

fwd-rule add vnid

Add external gateway forwarding rule.

Add a service port forwarding rule for an external gateway. This enables external networks to access a service hosted in the overlay network.

Syntax:

```
fwd-rule add vnid <VNID> overlay-ip <overlay IPv4>
floating-ip <floating IPv4 address> [proxy-min-ip <proxy start IP>]
[proxy-max-ip <proxy end IP>] [protocol <protocol>]
[port <logical port>] [overlay-port <overlay port>]
```

Parameters:

<VNID>	Range: 1-65535.																
<overlay IPv4>	IPv4 address for overlay service (VM).																
<floating IPv4 address>	Floating IPv4 address.																
<proxy start IP>	Specifies the first IPv4 address in the proxy IPv4 address range.																
<proxy end IP>	Specifies the last IPv4 address in the proxy IPv4 address range.																
<protocol>	Well-known protocol number (0-254). For example: <table><thead><tr><th>Number</th><th>Name</th></tr></thead><tbody><tr><td>0</td><td>Any (match any protocol)</td></tr><tr><td>1</td><td>CMP</td></tr><tr><td>2</td><td>IGMP</td></tr><tr><td>6</td><td>TCP</td></tr><tr><td>17</td><td>UDP</td></tr><tr><td>89</td><td>OSPF</td></tr><tr><td>112</td><td>VRRP</td></tr></tbody></table>	Number	Name	0	Any (match any protocol)	1	CMP	2	IGMP	6	TCP	17	UDP	89	OSPF	112	VRRP
Number	Name																
0	Any (match any protocol)																
1	CMP																
2	IGMP																
6	TCP																
17	UDP																
89	OSPF																
112	VRRP																
<logical port>	Port number for the gateway. Range: 1-65534.																
<overlay port>	Overlay service port (1-65535). This value is protocol dependent: 1-65535 for protocols 6 (TCP) and 17 (UDP) 0 for all other protocols.																

Command mode:

Service Gateway Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove-gateway)# fwd-rule add vnid 1
overlay-ip 192.168.1.2 floating-ip 20.20.20.1 protocol 6 port 5001 overlay-port 5001
```

fwd-rule add group_id

Add external gateway forwarding rule for a connectivity group.

Add a service port forwarding rule for an external gateway. This enables external networks to access a service hosted in the overlay network.

Syntax:

```
fwd-rule add group_id <Group ID> overlay-ip <overlay IPv4>
floating-ip <floating IPv4 address> [proxy-min-ip <proxy start IP>]
[proxy-max-ip <proxy end IP>] [protocol <protocol>]
[port <logical port>] [overlay-port <overlay port>]
```

Parameters:

<Group ID>	Connectivity group ID.																
<overlay IPv4>	IPv4 address for overlay service (VM).																
<floating IPv4 address>	Floating IPv4 address.																
<proxy start IP>	Specifies the first IPv4 address in the proxy IPv4 address range.																
<proxy end IP>	Specifies the last IPv4 address in the proxy IPv4 address range.																
<protocol>	Well-known protocol number (0-254). For example: <table><thead><tr><th>Number</th><th>Name</th></tr></thead><tbody><tr><td>0</td><td>Any (match any protocol)</td></tr><tr><td>1</td><td>ICMP</td></tr><tr><td>2</td><td>IGMP</td></tr><tr><td>6</td><td>TCP</td></tr><tr><td>17</td><td>UDP</td></tr><tr><td>89</td><td>OSPF</td></tr><tr><td>112</td><td>VRRP</td></tr></tbody></table>	Number	Name	0	Any (match any protocol)	1	ICMP	2	IGMP	6	TCP	17	UDP	89	OSPF	112	VRRP
Number	Name																
0	Any (match any protocol)																
1	ICMP																
2	IGMP																
6	TCP																
17	UDP																
89	OSPF																
112	VRRP																
<logical port>	Port number for the gateway. Range: 1-65534.																
<overlay port>	Overlay service port (1-65535). This value is protocol dependent: <table><tbody><tr><td>1-65535</td><td>for protocols 6 (TCP) and 17 (UDP)</td></tr><tr><td>0</td><td>for all other protocols.</td></tr></tbody></table>	1-65535	for protocols 6 (TCP) and 17 (UDP)	0	for all other protocols.												
1-65535	for protocols 6 (TCP) and 17 (UDP)																
0	for all other protocols.																

Command mode:

Service Gateway Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove-gateway)# fwd-rule add group_id 1
overlay-ip 192.168.1.2 floating-ip 20.20.20.1 protocol 6 port 5001 overlay-port 5001
```

fwd-rule delete vnid

Delete external gateway forwarding rule.

Syntax:

```
fwd-rule delete vnid <VNID> overlayip <overlay IPv4>
floating-ip <floating IPv4 address>
```

Parameters:

<VNID> Range: 1-65535.
<overlay IPv4> IPv4 address for overlay service (VM).
<floating IPv4 address> Floating IPv4 address.

Command mode:

Service Gateway Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove-gateway)# fwd-rule delete vnid 1
overlay-ip 192.168.1.2 floating-ip 20.20.20.1
```

fwd-rule delete group_id

Delete external gateway forwarding rule of a connectivity group.

Syntax:

```
fwd-rule delete group_id <Group ID> overlayip <overlay IPv4>
floating-ip <floating IPv4 address>
```

Parameters:

<Group ID> Connectivity group ID.
<overlay IPv4> IPv4 address for overlay service (VM).
<floating IPv4 address> Floating IPv4 address.

Command mode:

Service Gateway Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove-gateway)# fwd-rule delete group_id 1
overlay-ip 192.168.1.2 floating-ip 20.20.20.1
```

Miscellaneous Commands

clear logs

Deletes all log file contents.

Syntax:

clear logs

Command mode:

Global Configuration mode

Example:

```
SDN-VE @SDN-VE-Controller (config)# clear logs
```

exit

Exit from a context sub-mode and return to the parent mode. If already at the top level, exit from the command line interface and log out.

Syntax:

exit

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller (config-sdnve-dove-gateway)# exit
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# exit
SDN-VE @SDN-VE-Controller (config)# exit
SDN-VE @SDN-VE-Controller>
```

quit

Exit from the CLI and log out.

Syntax

quit

Command mode:

All

Example:

```
SDN-VE @SDN-VE-Controller (config)# quit
```

show

You can view SDN VE configuration and statistical information using a variety of `show` commands. For details, see [“Show Commands” on page 193](#).

Chapter 21. DSA Show Commands

Once you have logged in to a Distributed Services Appliance (DSA) module, you can view system configuration and statistical information using a variety of CLI `show` commands. The `show` commands are restricted from the User EXEC mode, but most are available globally in all other command modes.

This chapter discusses how to use each of the information-specific CLI `show` commands.

`show caroot-certificate`

Display CA root certificate information.

Syntax:

```
show caroot-certificate
```

Command mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show caroot-certificate
```

`show certificate`

Display DSA certificate information.

Syntax:

```
show certificate
```

Command mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show certificate
```

`show cli-timeout`

Any CLI session will be automatically logged out if idle for the length of time shown.

Syntax:

```
show cli-timeout
```

Command mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show cli-timeout
CLI TIMEOUT 5 min
```

See also:

`system cli timeout`

show config

Show the current DSA configuration properties. This command consolidates the following information from various other show commands.

- DSA Version
- IPv4 Management Configuration
- Service Appliance Configuration
- Certificate information

Syntax:

`show config`

Command mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show config
Software Version : 1.0.0.130603 Mon Jun  3 15:42:32 PDT 2013

ipmgmt set dhcp

dmc set ipv4 9.70.27.245 port 80
```

show crl

Display certificate revocation list (CRL) information.

Syntax:

`show crl`

Command mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show crl
```

show date

Displays system date.

Syntax:

show date

Command mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show date
2014-July-18
```

show dcs syslog

Show DCS system log messages.

Syntax:

show dcs syslog

Command mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show dcs syslog
DSA Version: 1.0.0.130603 Mon Jun  3 15:42:02 PDT 2013

DSA Version: 1.0.0.130603 Mon Jun  3 15:42:02 PDT 2013

DCS version 1.0.0.130603 Mon Jun  3 15:42:02 PDT 2013

Node IP 127.0.0.1, Port 902
Local DCS Service IP Address: 127.0.0.1
Node IP 127.0.0.1, Port 902
Existing Role File not found [/flash/dcs.role]
DCS: Local Node Inactive
DPS Protocol Handler Stopped
DPS Controller Interface Stopped
Wrote Role 0 to file /flash/dcs.role
DCS Role: Initialized to Inactive
DMC address has not been configured yet
getsockopt SO_RCVBUF returns 2001588984, rcv_size_len 0
setsockopt SO_RCVBUF set to 67108864
Adding Socket 35 to CORE API

DCS Server Started: IP Address <127.0.0.1>, Port <902>
Node IP 9.70.27.54, Port 902
----- Press any key to continue (q to quit) -----
```

show dgw syslog

Show DGW system log messages.

Syntax:

```
show dgw syslog
```

Command mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show dgw syslog
DOVE-GATEWAY: controller_interface/src/dgadmin_rest_client.c:dgwy_rest_client_t
o_dmc:884: Now send a RESTful HTTP request to Dove Controller, uri is /api/dove
/dgw/service/registration
.
.
.
----- Press any key to continue (q to quit) -----
```

show dmc-config

Show information about the DMC to which the DCS is connected.

Syntax:

```
show dmc-config
```

Command mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show dmc-config
DMC ipv4 : 9.70.27.245
DMC Port : 80
```

show dsa-version

Show Distributed Services Appliance software version information.

Syntax:

```
show dsa-version
```

Command mode:

Privileged EXEC and above

Example

```
SDN-VE-DSA# show dsa-version
Version: 1.2.0.140717 Thu Jul 17 00:51:33 PDT 2014
```

show ipmgmt

Shows DSA IPv4 address and netmask information.

Syntax:

```
show ipmgmt
```

Command mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show ipmgmt
Mgmt IPv4: 9.121.62.113
Mask: 255.255.254.0
Nexthop: 9.121.62.1
```

See also:

```
ipmgmt set ip
ipmgmt set cidr
ipmgmt set dhcp
```

show ipv4-interfaces

This command shows all IPv4 interfaces bound to the DSA. This information can be particularly important when configuring gateway (DGW) modules.

Syntax:

```
show ipv4-interfaces
```

Command Mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show ipv4-interfaces
0: 127.0.0.1
1: 9.70.27.54
```

show pvt-key

Display DSA private key information.

Syntax:

```
show pvt-key
```

Command Mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show pvt-key

-----BEGIN RSA PRIVATE KEY-----
MIIEpAIBAAKCAQEAXxiZU1WURAUv+jHVwIy1QmiQ1nwIi57UP053sv7bDf0LxJL
ayzw8BRHVxHyUD4i3KfWxPzsEuDIeae9TkpJpMpQcD9pOzT1GrHt/K+kf38+hkZC
iVKPvfpXsrkzuQUaqmW+Xw4UREW3Qdi6Qp57I7m9XL5X7XP43nVnEPnQuu4Sj0j4
6uhjJ6azAT9nakP1rsh4lN1+Jv8aEHikN1M9WbdTcs6ESnPI2FxadvJPAPderuhS
Lp4/IiMNrr03r7Y6ygG/6FFYf+50m8Bf7S8bH5/Fh/v9M/9XZUNdQwXNYR1PoP/G
hiZpXjRKC34EAnM3xdWmPG64Uxqx9/fa92Xq6wIDAQABAoIBAQCtmPs9qWTJ0bu4
PdCx7z5R8w9pEN7V6FGhRN5RjjU5GA4da1mf/bQst0vf/7z8wePEY2TzqnE02ZG1
prXPuYD4+sdBP7E6pE8ISEofHvM7VdC7aJDhDuTj0SwN8lIm3GLVWhIUk/JwXqKp
.
.
.
-----END RSA PRIVATE KEY-----
```

show security_mode

Display security configuration status.

Syntax:

```
show security_mode
```

Command Mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show security_mode

Authentication Disabled, CRL Disabled
```

show system acknowledgement

Show software licensing information for elements used in the DSA module.

Syntax:

```
show system acknowledgement
```

Command mode:

Privileged EXEC and above

Example:

```
SDN-VE-DSA# show system acknowledgement
```

show terminal-length

Show the number of lines displayed per screen. To facilitate reading lengthy output, the display for commands that produce more lines than defined by the terminal length will automatically pause, requiring a keypress before resuming each page of output.

Syntax:

```
show terminal-length
```

Command mode:

Privileged EXEC and above

Example:

```
DMC# show terminal-length
38 lines per screen
```

See also:

terminal-length

Chapter 22. DSA Configuration Commands

This chapter discusses how to use the individual CLI for making configuration changes.

Use the following command to access the configuration mode:

```
SDN-VE-DSA> enable  
SDN-VE-DSA# configure terminal  
SDN-VE-DSA(config)#
```

Clear Commands

clear screen

Clear the terminal screen.

Syntax:

```
clear screen
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE-DSA(config)# clear screen
```

clear gwstats

Clear the gateway statistics that are collected by a DSA operating as a DGW node.

Syntax:

```
clear gwstats
```

Command mode:

Global Configuration mode

Example:

```
SDN-VE-DSA(config)# clear gwstats
```

CLI Timeout Commands

cli timeout

Sets length of time before CLI times out.

Syntax:

```
cli timeout mins <minutes>
```

Parameters:

<minutes> Timeout period in minutes. Value: 1-60.

Command mode:

Global Configuration mode

Example:

```
SDN-VE-DSA(config)# cli timeout 8
```

See also:

```
show cli-timeout
```

Controller Commands

dmc set ip

Bind DSA to DMC.

Define the IPv4 address of the DMC to which this DSA will register.

Syntax:

```
dmc set ip addr <IPv4 address>
```

Parameters:

<IPv4 address> IPv4 address in dotted decimal notation (a.b.c.d)

Command mode:

Global Configuration mode

Example:

SDN-VE-DSA(config)# dmc set ip addr 10.10.0.10
--

Hostname Commands

hostname set

Set DSA hostname.

Syntax:

```
hostname set name
```

Command mode:

Global Configuration mode

hostname reset

Reset the DSA hostname.

Syntax:

```
hostname reset
```

Command mode:

Global Configuration mode

Image Upgrade Commands

dsa-upgrade

Upgrade the DSA software image. The new image file must be accessible to the DSA. Once upgraded, the DSA will automatically reboot in order to run the new image. You can verify the upgrade by using the `show dsa-version` command.

Syntax:

```
dsa-upgrade url <image URL>
```

Parameters:

<image URL> URL (1 to 128 characters) for the DSA software image file.

Command mode:

Global Configuration mode

Example:

```
SDN-VE-DSA(config)# system dsa-upgrade url
ftp://9.111.86.13/xyz/ibm-sdn-dsa-upgrade-1.0.0.img
Please wait while the DSA Image is being upgraded!!
SDN-VE-DSA(config)#
```

See also:

```
show dsa-version
reload
```

IP Management Commands

ipmgmt set

Set a static DSA address by specifying an IPv4 address and netmask or CIDR designation.

Syntax:

```
ipmgmt set ip addr <IPv4 address> mask <netmask>
ipmgmt set cidr addr <CIDR>
```

Parameters:

<IPv4 address>	IPv4 address in dotted decimal notation (a.b.c.d)
<netmask>	IPv4 netmask in dotted decimal notation (a.b.c.d)
<CIDR>	IPv4 address in CIDR format (a.b.c.d/e)

Command mode:

Global Configuration mode

Examples:

SDN-VE-DSA(config)# ipmgmt set ip addr 10.10.0.10 mask 255.255.255.0
--

SDN-VE-DSA(config)# ipmgmt set cidr 10.10.0.10/24

See also:

```
show ipmgmt
ipmgmt set dhcp
```

ipmgmt set dhcp

Set a dynamic DSA address via DHCP.

Note: Setting the IPv4 address to use DHCP clears the static DSA IPv4 address and gateway.

Syntax:

```
ipmgmt set dhcp
```

Command mode:

Global Configuration mode

Example:

SDN-VE-DSA(config)# ipmgmt set dhcp

See also:

```
show ipmgmt
ipmgmt set
```

ipmgmt set nexthop

Set the DSA gateway address by specifying a static IPv4 address and netmask.

Syntax:

```
ipmgmt set nexthop ip <IPv4 address>
```

Parameters:

<IPv4 address> IPv4 address in dotted decimal notation (a.b.c.d)

Command mode:

Global Configuration mode

Example:

```
SDN-VE-DSA(config)# ipmgmt set nexthop ip 10.10.0.1
```

Password Configuration Commands

password

Change the DSA administrator password.

The password length must be at least 6 and no longer than 31 characters.

Syntax:

password

Command mode:

Global Configuration mode

Example:

```
SDN-VE-DSA(config)# password
Enter new admin password: *****
Verify new admin password: *****
Success: admin password changed!
```

Miscellaneous Commands

exit

Exit from a context sub-mode and return to the parent mode. If already at the top level, exit from the command line interface and log out.

Syntax:

exit

Command mode:

All

Example:

```
SDN-VE-DSA(config)# exit
SDN-VE-DSA# exit
SDN-VE-DSA>
```

See also:

quit

find

Find Command Syntax.

List the commands available in the current mode.

Syntax:

find [keyword]

Command mode:

All

ping

Use the ping utility to test network connectivity.

Syntax:

ping dst <destination IPv4> [src <source IPv4>]

Parameters:

<destination IPv4> Destination IPv4 address.

<source IPv4> Optional source IPv4 address.

Command mode:

All

Example:

```
SDN-VE-DSA# ping dst 9.0.130.50
PING 9.0.130.50 (9.0.130.50): 56 data bytes
64 bytes from 9.0.130.50: seq=0 ttl=115 time=76.719 ms
64 bytes from 9.0.130.50: seq=1 ttl=115 time=76.528 ms
64 bytes from 9.0.130.50: seq=2 ttl=115 time=76.730 ms
64 bytes from 9.0.130.50: seq=3 ttl=115 time=78.763 ms
--- 9.0.130.50 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max = 76.528/77.185/78.763 ms
```

quit

Exit from the CLI and log out.

Syntax

`quit`

Command mode:

All

Example:

```
SDN-VE-DSA(config)# quit
```

See also:

`exit`

reload

Reset and reboot the DSA module.

Syntax:

`reload`

Command mode:

Global Configuration mode

Example:

```
SDN-VE-DSA(config)# reload
Reload DMC [y|n]?: y
Please wait while the DSA is being reloaded!!
```

show

You can view DSA configuration and statistical information using a variety of `show` commands. For details, see [“DSA Show Commands” on page 289](#).

PKI Configuration Commands

security-mode auth

Enable or disable SSL authentication security mode. You must first import the certificates before enabling authentication. See [“ssl-import ca-root-cert” on page 308](#), [“ssl-import certificate” on page 309](#), [“ssl-import crl” on page 309](#), [“ssl-import key” on page 310](#).

For details on the authentication feature, see [Chapter 12, “Public Key Infrastructure”](#).

Syntax:

```
security-mode auth set {enable|disable}
```

Command mode:

Global Configuration mode.

Example:

```
SDN-VE-DSA (config)# security-mode auth set enable
```

security-mode crl

Enable or disable CRL security mode.

CRL (Certificate Revocation List) feature verifies the certificate status for TLS/IPSec/DTLS sessions against an uploaded CRL file. If CRL verification is enabled, session handshake will fail if either side presents a certificate that has been revoked for any reason.

Syntax:

```
security-mode acrl set {enable|disable}
```

Command mode:

Global Configuration mode.

Example:

```
SDN-VE-DSA (config)# security-mode crl set enable
```

ssl-import ca-root-cert

Import Certificate Authority (CA) root certificate.

Syntax:

```
ssl-import ca-root-cert url <URL>
```

Parameters:

<URL> Location of the certificate. Length: 1-128.

Command mode:
Global Configuration mode.

Example:

```
SDN-VE-DSA (config)# ssl-import ca-root-cert url ftp://9.0.130.50/cert/cacert.pem
```

ssl-import certificate

Import DSA certificate.

Syntax:

```
ssl-import certificate url <URL>
```

Parameters:

<URL> Location of the certificate. Length: 1-128.

Command mode:
Global Configuration mode.

Example:

```
SDN-VE-DSA (config)# ssl-import certificate url ftp://9.0.130.50/cert/dsacert.pem
```

ssl-import crl

Import certificate revocation list.

Syntax:

```
ssl-import crl url <URL>
```

Parameters:

<URL> Location of the certificate. Length: 1-128.

Command mode:
Global Configuration mode.

Example:

```
SDN-VE-DSA (config)# ssl-import crl url ftp://9.0.130.50/cert/crl.pem
```

ssl-import key

Import DSA private key.

Syntax:

```
ssl-import key url <URL>
```

Parameters:

<URL> Location of the certificate. Length: 1-128.

Command mode:

Global Configuration mode.

Example:

```
SDN-VE-DSA (config)# ssl-import key url ftp://9.0.130.50/cert/key.pem
```

Terminal Length Configuration Commands

terminal-length

Set the number of lines available on the terminal display. To facilitate reading lengthy output, the display for commands that produce more lines than defined by the terminal length will automatically pause, requiring a keypress before resuming each screen of output.

Syntax:

```
terminal-length length <lines>
```

Parameters:

<lines>	Number of lines per screen (1-256), or 0 to permit unlimited lines per screen.
---------	--

Command mode:

Global Configuration mode

Example:

```
SDN-VE-DSA(config)# terminal-length length 24
```

See also:

```
show terminal-length
```

Chapter 23. Diagnostics Commands

Diagnostics commands can be viewed in diagnostic context.

To enter the diagnostics context mode:

```
SDN-VE@SDN-VE-Controller> diagnostic terminal
SDN-VE@SDN-VE-Controller (diagnostic)# ?
copy          Saving configurations
exit          Exits from diagnostic mode
history       Displays current session's command line history
ping          Ping a IP address
quit          Aborts the CLI session
show          Shows configuration
system        Run System commands
tcpdump       Run tcpdump to monitor packets on a network interface
traceroute    Trace an IP address
```

ping

Pings an IP Address

Syntax:

`ping <Target IPv4 address>`

Command mode:

Diagnostic Context

Example:

```
SDN-VE @SDN-VE-Controller (diagnostic)# ping 9.121.62.23
64 bytes from 9.121.62.23: seq=0 ttl=64 time=2.363 ms
64 bytes from 9.121.62.23: seq=1 ttl=64 time=0.261 ms
64 bytes from 9.121.62.23: seq=2 ttl=64 time=0.343 ms

--- 9.121.62.23 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.261/0.989/2.363 ms
```

tcpdump

Run tcpdump to monitor packets on a network interface

Syntax:

```
tcpdump device <interface name> [<options>]
```

Parameter:

<interface name>	Name of the interface
<options>	Select from the following options to restrict the output:
time	Time in seconds to run the command. (default=10, if the number of packets are not specified)
port	Port number to monitor
prot	Protocol name
not-port	Port number to be excluded from monitoring
packets	Number of packets to be captured
file	Dump packets to a file
src-host	Packets from a source host
src-net	Packets from a source network
dest-host	Packets to a destination host
dest-net	Packets to a destination network

Command mode:

Diagnostic Context

Example:

```
SDN-VE @SDN-VE-Controller (diagnostic)# tcpdump device eth0
(Type 'q' at any time to quit)
```

Note: tcpdump file can be accessed from the browser interface using
<http://<controller IP address>/log/plog/>

Logs are also available at: <http://<controller IP address>/log>

traceroute

Traces the target IP address

Syntax:

`traceroute <IPv4 address>`

Parameter :

`<IPv4 address>` Target IPv4 Address

Command mode:

Diagnostic Context

Example:

```
SDN-VE @SDN-VE-Controller (diagnostic)# traceroute 9.121.62.23
traceroute to 9.121.62.23 (9.121.62.23), 30 hops max, 46 byte packets
 1  9.121.62.23  4.245 ms
```


Part 4: Appendices

Appendix A. New and Updated Features

IBM SDN VE 1.2 has also been updated to include several new/updated features, summarized in the following sections.

CLI

The IBM SDN VE Controller command-line interface (CLI) has changed significantly. The updated commands are not executable in releases prior to 1.2.

Following is a summary of the CLI changes. See [Part 3, “Command Reference,” on page 185](#) for details.

DOVECLI

The DOVE CLI, which was used to access the SDN-VE context, has been removed.

Old commands

```
SDN-VE @SDN-VE-Controller > configure terminal
SDN-VE @SDN-VE-Controller (config)# sdnve-dove terminal
SDN-VE @SDN-VE-Controller (config-sdnve-dove)# dovecli
SDN-VE-Controller>
```

In this release, all SDN VE DOVE configuration is performed using the SDN VE DOVE CLI:

```
SDN-VE @SDN-VE-Controller > configure terminal
SDN-VE @SDN-VE-Controller (config)# sdnve-dove terminal
SDN-VE-Controller (config-sdnve-dove)#
```

The SDN VE DOVE configuration includes the following:

- External IP address of the controller HA nodes (See: [“external-ip” on page 274](#))
- HA start or stop (See: [“ha start” on page 275](#) and [“ha stop” on page 275](#))
- HA Synchronization (See: [“ha-synchronization start” on page 275](#))
- HA peers (See: [“peers” on page 276](#))
- Service appliance configuration: DGW/DCS role setting; External and VLAN gateway configuration (See: [“service dgw” on page 277](#); [“service role dcs” on page 278](#); [“service role dgw” on page 278](#); [“Service Gateway Configuration Mode Commands” on page 284](#))
- Underlay networks (See: [“underlay-network subnet” on page 281](#))
- Virtual Router Redundancy Protocol (VRRP) (See: [“vrrp add” on page 282](#))

High-Availability (HA)

Configuring IBM SDN VE controller nodes HA is now a two-step process. You must first configure the SDN-VE HA (See [“Establish SDN VE Controller HA” on page 43](#)) and then the SDN VE DOVE HA (See [“Configure SDN VE DOVE HA” on page 45](#)).

NIST

The IBM SDN VE implementation is compliant with National Institute of Standards and Technology (NIST) Special Publication (SP) 800-131A.

The IBM SDN VE controller can be configured to operate in NIST-compliant mode. By default, NIST is disabled on the controller.

See [Chapter 11, “NIST”](#) for details.

Overlay Networks

In releases prior to 1.2, overlay networks were defined with domains, virtual networks, the address spaces mapped to the networks, and the policies between networks.

In this release, the overlay network consists of tenants, connectivity groups, the address spaces mapped to the connectivity groups, and the policies between connectivity groups.

PKI

A Public Key Infrastructure (PKI) assures secure exchange of data using a public and a private cryptographic key pair. This key pair is exchanged via a trusted authority.

PKI includes the following:

- Certificate authority (CA): Issues and verifies digital certificates.
- Registration authority (RA): Verifies identity of the users/applications that request information from the CA.

The IBM SDN VE Controller and the Distributed Service Appliance (DSA) can be configured to use PKI. By default, security is enabled and authentication is disabled.

See [Chapter 12, “Public Key Infrastructure”](#) for details.

QoS

Quality of Service (QoS) features allow you to allocate network resources to mission-critical applications at the expense of applications that are less sensitive to such factors as time delays or network congestion. You can configure your network to prioritize specific types of traffic, ensuring that each type receives the appropriate QoS level.

See [Chapter 14, “Quality of Service”](#) for details.

TSO

TCP Segmentation Offload (TSO) breaks down large groups of data (TCP packets) sent over a network into smaller segments. TSO improves network performance by reducing the CPU overhead.

See [Chapter 16, “TCP Segmentation Offload”](#) for details.

VRRP

The IBM SDN VE solution supports IPv4 high-availability (HA) network topologies through implementation of the Virtual Router Redundancy Protocol (VRRP). VRRP enables redundant router configurations within a LAN, providing alternate router paths for a host to eliminate single points-of-failure within a network.

See [Chapter 17, “Virtual Router Redundancy Protocol”](#) for details.

Waypoint Connectivity Service

A middlebox is a network appliance that resides between the source and destination of a packet. Typical middlebox examples include firewalls, Network Address Translators (NAT), load balancers, and Intrusion Prevention Systems (IPS)/Intrusion Detection Systems (IDS).

The IBM SDN VE solution supports routing of traffic through such middleboxes. Use of middleboxes for routing is also known as service insertion or Waypoint service enablement.

See [Chapter 10, “Waypoint Connectivity Service”](#) for details.

Appendix B. OpenStack Neutron APIs

The table listed below captures the indicative list of the OpenStack Neutron APIs. For more details on each API, please visit the OpenStack website: Networking v2.0 APIs.

Table 24. OpenStack Neutron APIs

API	Purpose	Operation
Networks		
/	Lists information about all Networking API versions.	GET
/v2.0	Shows details for Networking API v2.0.	GET
/v2.0/extensions	Lists available Networking API extensions.	GET
/v2.0/extensions/{alias}	Gets detailed information for a specified extension.	GET
/v2.0/networks	Lists networks to which the specified tenant has access.	GET
/v2.0/networks	Creates a network.	POST
/v2.0/networks	Creates multiple networks in a single request.	POST
/v2.0/networks/{network_id}	Shows information for a specified network.	GET
/v2.0/networks/{network_id}	Updates a specified network.	PUT
/v2.0/networks/{network_id}	Deletes a specified network and its associated resources.	DELETE
Subnets		
/v2.0/subnets	Lists subnets to which the specified tenant has access.	GET
/v2.0/subnets	Creates a subnet on a specified network.	POST
/v2.0/subnets	Creates multiple subnets in a single request. Specify a list of subnets in the request body.	POST
/v2.0/subnets/{subnet_id}	Shows information for a specified subnet.	GET
/v2.0/subnets/{subnet_id}	Updates a specified subnet.	PUT
/v2.0/subnets/{subnet_id}	Deletes a specified subnet.	DELETE
Ports		
/v2.0/ports	Lists ports to which the tenant has access.	GET
/v2.0/ports	Creates a port on a specified network.	POST
/v2.0/ports	Creates multiple ports in a single request. Specify a list of ports in the request body.	POST

Table 24. OpenStack Neutron APIs

API	Purpose	Operation
/v2.0/ports/{port_id}	Shows information for a specified port.	GET
/v2.0/ports/{port_id}	Updates a specified port.	PUT
/v2.0/ports/{port_id}	Deletes a specified port.	DELETE
Extensions		
/v2.0/extensions	Lists available Networking API extensions.	GET
/ports	Lists ports to which the tenant has access.	GET
/ports/{port-id}	Shows information for a specified port.	GET
/ports	Creates a port on a specified network.	POST
/ports/{port-id}	Updates a specified port.	PUT
(See: The binding Extended Attributes for Ports)		
/v2.0/routers	Creates a logical router.	POST
/v2.0/routers/{router_id}	Shows details for a specified router.	GET
/v2.0/routers/{router_id}	Updates a logical router.	PUT
/v2.0/routers/{router_id}	Deletes a logical router. Also deletes its external gateway interface, if present.	DELETE
/v2.0/routers/{router_id}/add_router_interface	Adds an internal interface to a logical router.	PUT
/v2.0/routers/remove_router_interface	Removes an internal interface from a logical router.	PUT
/v2.0/floatingips/{floatingip_id}	Shows details for a specified floating IP.	GET
/v2.0/floatingips	Creates a floating IP. If port information is specified, associates the floating IP with an internal port.	POST
/v2.0/floatingips/{floatingip_id}	Updates a floating IP and its association with an internal port.	PUT
/v2.0/floatingips/{floatingip_id}	Deletes a floating IP. Also deletes its associated port, if present.	DELETE

The following additional General APIs are also supported:

- **Authentication and Authorization:** Specify access criteria based on operations or resource.
- **Filtering and column selection:** Filtering based on all top level attributes of a resource. Filters are applicable to all list requests.
- **Bulk create:** Create several objects of the same type in a single API request.

The `binding` Extended Attributes for Ports

The attributes can be used with the APIs to get more information about ports, and to create and update port objects.

Following are the `binding` Extended Attributes for Ports:

- `binding:vif_type` - R¹
- `binding:host_id` - CRU
- `binding:profile` - CRU
- `binding:capabilities` - R

-
- 1. C. Used in create operations.
 - R. This attribute is returned in response to show and list operations.
 - U. The value of this attribute can be updated.
 - D. The value of this attribute can be deleted.

Appendix C. REST API

The table listed below captures the indicative list of the REST API details. Nevertheless, for the comprehensive list along with details like data model and operations, please refer to REST APIs file packaged along with the software files.

The REST APIs can be accessed using the following format:

`https://<Controller HA external IPv4 address>:8443/<Module Prefix><REST API>`

For Example:

`https://9.121.84.20:8443/controller/nb/v2/flowgroupmanager/flowgroup?filter=dynamic`

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
FlowGroupsManager: /controller/nb/v2/flowgroupmanager				
/flowgroup	Get all flow groups.	GET	Filter/query	404
/flowgroup?filter=dynamic	Get All Dynamic Flow-Groups.			
/flowgroup?filter=static	Get All Static FlowGroups.			
/flowgroup	Add flow groups.	POST	admin/query	503; 500
/flowgroup?admin=TRUE	Add static flow group.			
/flowgroup?admin=FALSE	Add non-admin static flow group (Used by non-CLI NBI client), default if admin value not specified.			
/flowgroup	Delete a flow group.	DELETE		503; 500
/flow/{groupname}	Add static flow in group.	POST	groupname/path	404; 405; 406
/flow/{groupname}	Modify static flow in group.	PUT	groupname/path	404; 405; 406
/flowgroup/{groupname}	Get flow group.	GET	groupname/path	404
/flowgroup/{groupname}?filter=static	Get static flow group.	GET		
/flowgroup/{groupname}?filter=dynamic	Get dynamic flow group.	GET		
/flowgroup/{groupname}	Modify static flow group.	PUT	groupname/path	503; 500
/flowgroup/{groupname}	Remove flow group.	DELETE	groupname/path	503; 500
/flow/{groupname}/{flowname}	Delete flow in group.	DELETE	groupname/path flowname/path	404; 405
/flowgroup/{groupname}/install	Install the FlowGroup.	POST	groupname/path	503; 500

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
/flowgroup/{groupname}/install	Uninstall the FlowGroup.	DELETE	groupname/path	503; 500
/switchid/{switchid}/flow	Get All Flows for the Switch.	GET	switchid/path filter/query	404
/switchid/{switchid}/flow? filter=dynamic	Get Dynamic Flows.	GET		
/switchid/{switchid}/flow? filter=static	Get Static Flows.	GET		
InterfaceManager: /controller/nb/v2/interfacemanager				
/interface	Create interface.	POST		503; 415; 409
/interfaces	Get list of interfaces.	GET	tenantid/query networkad- dress/query datalinkad- dress/query vlan/query vnid/query	503; 415
/interfaces?tenantid={tenantId}		GET		
/interfaces? networkaddress={networkAddress}		GET		
/interfaces? datalinkaddress={datalinkAddress}		GET		
/interfaces?vlan={vlanId}		GET		
/interfaces?vnid={vnid}		GET		
/interface/l3	Create interface.	POST		503; 415; 409
/hostnodeconnector/ datalinkaddress/{datalinkAddress}	Get host node connector.	GET	datalinkaddress/path	503; 415
/hostnodeconnector/ interfaceid/{interfaceId}	Get host node connector.	GET	interfaceid/path	503
/interface/interfaceid/{interfaceId}	Get interface.	GET	interfaceid/path	503; 415
/interface/interfaceid/{interfaceId}	Update interface.	PUT	interfaceid/path	503
/interface/interfaceid/{interfaceId}	Delete interface.	DELETE	interfaceid/path	503
/interface/datalinkaddress/{data- linkAddress}/networkaddress/{net- workAddress}	Get interface.	GET	datalinkaddress/path networkad- dress/path	503; 415
/hostnodeconnector/tenantid/{tenant- Id}/datalinkaddress/{datalinkAd- dress}/networkaddress/{networkAd- dress}	Get host node connector.	GET	tenantid/query datalinkaddress/path networkad- dress/path	503; 415

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
RouteManager: /controller/nb/v2/l3				
/l3route	Get all routes.	GET	tenant/query ipv4address/query ipv4gateway/query	404
/l3route?tenant=x	Get all Route from routing table for given tenant-id.	GET		
/l3route?ipv4address=x	Read route table routes for given IP Address (IPv4).	GET		
/l3route?ipv4gateway=x	Read route table routes for given gateway address (next hop).	GET		
/l3router	Get all L3 Routers for given tenant-id.	GET	tenant/query	404
/l3router	Create L3 Router for given tenant-id.	POST	tenant/query	404
/l3router/summary	Get Route Summary for given tenant-id.	GET	tenant/query	404
/l3route/routerid/{routerid}	Get all L3 Routers for given tenant-id or router-id.	GET	tenant/query routerid/query	404
/l3router/routerid/{routerid}	Get all L3 Routers for given tenant-id or router-id.	GET	tenant/query routerid/pathy	404
/l3router/routerid/{routerid}	Update L3 Routers for given tenant-id or router-id.	PUT	tenant/query routerid/path ipv4gateway/query	404
/l3router/routerid/{routerid}	Delete L3 Routers for given tenant-id or router-id.	DELETE	tenant/query routerid/path	404
/l3router/routerName/{name}	Create a L3-V-Router for the tenant.	POST	tenant/query name/path	404
/l3router/routerName/{name}	Delete a L3-V-Router for the tenant.	DELETE	tenant/query name/path	404
/l3route/ip/routerid/{routerid}	Get IP information of L3 route.	GET	tenant/query ipv4address/query ipv4gateway/query routerid/path	404
/l3router/routerId/{routerId}/l3interface	Create V-Router Interface.	POST	tenant/query routerid/path	404
/l3router/routerId/{routerId}/l3route	Create L3 route.	POST	tenant/query routerid/path	404
/l3router/routerId/{routerId}/summary	Get summary of L3 routes.	GET	tenant/query routerid/path	404

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
/l3router/routerId/{routerId}/ipv4gateway/{ipv4gateway}	Delete IPv4 gateway.	DELETE	tenant/query ipv4gateway/query routerid/path	404
/l3router/routerId/{routerId}/l3interface/{l3interfaceId}	Update interface.	PUT	tenant/query routerid/path interfaceid/path	404
/l3router/routerId/{routerId}/l3interface/{l3interfaceId}	Delete interface.	DELETE	tenant/query routerid/path interfaceid/path	404
/l3router/routerId/{routerId}/l3route/{routeId}	Update route ID.	PUT	tenant/query routeid/query routerin/query	404
/l3router/routerId/{routerId}/l3route/{routeId}	Delete route ID.	DELETE	tenant/query routeid/query routerin/query	404
StaticRouting: /one/nb/v2/l3/				
/l3/{containerName}	Get L3 information.	GET	containername/path	404
/l3/{containerName}	Create L3 container.	POST	containername/path	404
LogicalGroups: /controller/nb/v2/l3n/				
/endPoints	List all endpoints.	GET	group/query subnet/query tenant/query domain_type/query	404
/endPoints	Create an endpoint with endpoint attributes.	POST		409; 404
/endPoints/{id}	Get endpoint information for a specified group ID.	GET	id/path	404
/endPoints/{id}	Update endpoint information for a specified group ID.	PUT	id/path	404
/endPoints/{id}	Delete endpoint for a specified group ID.	DELETE	id/path	404
/groups	List all groups.	GET	tenant/query domain_type/query waypoint/query	404
/groups	Create a group with group attributes.	POST		404
/groups/connectivity	Get information of a connectivity policy.	GET	tenant/query service/query	404
/groups/connectivity	Create a connectivity policy with specified attributes.	POST	tenant/query	404

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
/groups/vnidexport	Export a specified group.	POST		408
/groups/{id}	Get information of a group for a specified ID.	GET	id/path	404
/groups/{id}	Update a group with specified group attributes.	PUT	id/path	404
/groups/{id}	Delete a group.	DELETE	id/path	404
/groups/connectivity/{policy_id}	Get the connectivity policy for the given ID.	GET	policy_id/path	404
/groups/connectivity/{policy_id}	Delete a connectivity policy.	DELETE	tenant/query policy_id/path	404
/groups/connectivity/adddnat/{policy_id}	Updates a connectivity policy with DNAT rules. (Applicable only to DOVE networks.)	PUT	tenant/query policy_id/path	404
/groups/connectivity/removednat/{policy_id}	Updates a connectivity policy by removing DNAT rules. (Applicable only to DOVE networks.)	PUT	tenant/query policy_id/path	404
/groups/connectivity/{grp1_id}/{grp2_id}	Get the connectivity policy for the specified ID.	GET	tenant/query grp1_id/path grp2_id/path	404
/groups/connectivity/{grp1_id}/{grp2_id}	Updates a connectivity policy.	PUT	tenant/query grp1_id/path grp2_id/path	404
/groups/connectivity/{grp1_id}/{grp2_id}	Deletes a connectivity policy.	DELETE	tenant/query grp1_id/path grp2_id/path	404
/groups/vnidexport/{group_id}/{ip_addr}	Unexport a group.	DELETE	group_id/path ip_addr/path	404
/groups/{id}/add_subnet/{subnet}	Add a subnet to a group.	PUT	id/path subnet/path	404
/groups/{id}/delete_subnet/{subnet}	Delete a subnet from a group.	PUT	id/path subnet/path	404
/networks	List all networks.	GET	tenant/query domain_type/query	404
/networks	Create a network with network attributes.	POST	tenant/query domain_type/query	404
/networks/{id}	Get network information.	GET	id/path	404
/networks/{id}	Update a network with specified attributes.	PUT	id/path	404

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
/networks/{id}	Delete a network.	DELETE	id/path	404
/ports	List all ports.	GET	group/query subnet/query tenant/query mac_address/query ip_address/query domain_type/query	404
/ports	Create a port with port attributes.	POST		409; 404
/ports/{id}	Get information about a port.	GET	id/path	404
/ports/{id}	Update port information.	PUT	id/path	404
/ports/{id}	Delete a port.	DELETE	id/path	404
/routers	List all routers.	GET	tenant/query domain_type/query	404
/routers	Create a router with router attributes.	POST		404
/routers/{id}	Get information about a router.	GET	id/path	404
/routers/{id}	Update router information.	PUT	id/path	404
/routers/{id}	Delete a router.	DELETE	id/path	404
/routers/{id}/add_router_interface	Add an interface to a router.	PUT	id/path	404
/routers/{id}/remove_router_interface	Delete a router interface.	PUT	id/path	404
/subnets	List all subnets.	GET	group/query tenant/query domain_type/query	404
/subnets	Create a subnet with subnet attributes.	POST		409; 404
/subnets/{id}	Get information about a subnet.	GET	id/path	404
/subnets/{id}	Update subnet information.	PUT	id/path	404
/subnets/{id}	Delete a subnet.	DELETE	id/path	404
/tenants	List all tenants.	GET	domain_type/query	404
/tenants	Create a tenant with tenant attributes.	POST		404
/tenants/{id}	Get information about a tenant.	GET	id/path	404

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
/tenants/{id}	Update tenant information.	PUT	id/path	404
/tenants/{id}	Delete a tenant.	DELETE	id/path	404
PolicyManager: /one/nb/v2/policymgr				
/policies	List all ACL policies.	GET	tenant/query	404
/policies	Create a policy with the specified attributes.	POST		404; 400
/policysetinstalls	Installs a policy set on a specified target.	POST		404; 400
/policysetinstalls	List all the policy sets to be installed.	GET	tenant/query	404
/policysets	List all the policy sets.	GET	tenant/query	404
/policysets	Create a policy set with the specified attributes.	POST		404; 409
/policies/mb	List all the MiddleBox policies.	GET	tenant/query	404
/policies/mb	Create a MiddleBox policy with the specified attributes.	POST		404; 400
/policies/mb/{id}	Get the policy for the specified ID.	GET	id/path tenant/query	404
/policies/{id}	Get the policy of the specified ID.	GET	id/path tenant/query	404
/policies/{id}	Modify the policy of the specified ID.	PUT	id/path tenant/query	404
/policies/{id}	Delete a policy.	DELETE	id/path tenant/query	404
/policysetinstalls/{id}	List all policy to be installed for a specified ID.	GET	id/path tenant/query	404
/policysetinstalls/{id}	Delete all installed policies for a specified ID.	DELETE	id/path tenant/query	404
/policysets/{id}	Get the policy set of the specified ID.	GET	id/path tenant/query	404
/policysets/{id}	Modify the policy set of the specified ID.	PUT	id/path tenant/query	404; 409
/policysets/{id}	Delete a policy set.	DELETE	id/path tenant/query	404
/policysets/{id}/policy	Modifies a policy set of the specified ID by adding the policy to the set.	PUT	id/path tenant/query	400; 404; 409

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
/policysets/{id}/policy/{policyid}	Modifies a policy set of the specified ID by deleting a policy from the set.	DELETE	id/path tenant/query policyid/path	404; 409
Replication: /controller/nb/v2/replication				
/fr	Create a replication/redirection session,	POST		403; 415; 500; 503
/fr	Get information about all replication/redirection sessions.	GET	mode/query tenant/query	403; 503
/fr	Modify all replication/redirection sessions.	PUT		403; 415; 503
/fr	Delete all replication/redirection sessions.	DELETE	mode/query	403; 415; 503
/fr/{sessionName}	Get information about a specific replication/redirection session.	GET	sessionName/path mode/query tenant/query	403; 503
/fr/{sessionName}	Delete a specific replication/redirection session.	DELETE	sessionName/path mode/query	403; 503
/fr/{sessionName}/start	Start a specific replication/redirection session.	PUT	sessionName/path	403; 503
/fr/{sessionName}/stop	Stop a specific replication/redirection session.	PUT	sessionName/path	403; 503
/fr/{sessionName}/start/{mode}	Start a replication/redirection session for the specified mode.	PUT	sessionName/path mode/query	403; 503
/fr/{sessionName}/stop/{mode}	Stop a replication/redirection session for the specified mode.	PUT	sessionName/path mode/query	403; 503
SPARTA: /controller/nb/v2/sparta				
/destTree/mac/{hostMac}	Get the SPARTA Tree for the specified host.	GET	hostMac/path	200; 404; 500; 503
/path/srcMac/{srcHost}/destMac/{dstHost}	Get the SPARTA Tree for the specified host.	GET	srcHost/path dstHost/path	200; 404; 500; 503
UserManager: : /controller/nb/v2/usermanager				
/users	Get a list of users.	GET	tenant/query	503
/users	Create a user.	POST		400; 403; 503
/users	Update a user.	PUT		400; 403; 503
/user/current	Get details on the user currently logged in.	GET		503

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
/user/isAuthorized	Get information about the current user's authorization; if the user is authorized to continue with the session.	GET		503
/user/password	Update user password.	PUT		400; 403; 503
/user/registerLogin	Provides access if the user has successfully registered a session.	POST		503
/user/registerLogoff	Provides access if the user has successfully registered a session.	POST		503
/users/active	Provides a list of all logged in users; third-party authenticated, local user authenticated.	GET		503
/user/password/reset	Update user password.	PUT		400; 403; 503
/user/sessionInactiveTimeout/{timeoutInterval}	Set the session inactive timeout (in seconds) for the current user session.	POST		503
/users/user/{userName}	Get information about the user currently logged in.	GET	userName/path	503
/users/user/{userName}	Delete a user.	DELETE	userName/path	503
Clustering: /controller/nb/v2/clustering				
/ha/cluster	Displays the cluster info i.e. cluster name & members details	GET		503
/ha/disconnect	Disconnects the specific Node from the cluster	PUT		404
/ha/rejoin	Rejoins the Node to the same cluster when it is disconnected	PUT	containerName/path	404
/ha/cluster/name/{clusterName}	Helps in forming the cluster as well as updating the cluster members	PUT	clusterName/path	404
Multicast: /controller/nb/v2/multicast				
/groups	Get multicast information of the group.	GET	tenant/query multicastip/query	200; 404; 503
/groups	Delete multicast IP of the group.	DELETE	tenant/query multicastip/query	200; 404; 500; 503
/vlanrange	Get information about the VLANs.	GET		200; 404; 503

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
/vlanrange	Modify VLANs for the given range.	PUT	lower/query upper/query	201; 404; 503
/vlanrange	Delete a range of VLANs.	DELETE		200; 404; 500; 503
/querier/interval	Get querier interval information.	GET		200; 404; 503
/querier/interval	Update querier interval.	PUT	interval/query	201; 404; 503
/tree/sender/{srcMac}	Delete source MAC address.	DELETE	tenant/query multicastip/query srcMac/path	200; 404; 500; 503
/tree/sender/{srcMac}	Get source MAC information.	GET	tenant/query multicastip/query srcMac/path	200; 404; 503
Statistics: /controller/nb/v2/statistics				
/flowstats	Get all flow statistics	GET	switch/query flow/query	200; 404; 503
/hoststats	Get all host statistics.	GET	addr/query	200; 404; 503
/portstats	Get all port statistics.	GET	port/query node/query	200; 404; 503
/switchstats	Get all switch statistics.	GET	switch/query	200; 404; 503
/tenantstats	Get statistics of the specified tenant.	GET	tenant/query	200; 404; 503
/tenantstats/all	Get all tenant statistics.	GET	tenant/query	200; 404; 503
Waypoint: /controller/nb/v2/waypoint				
MiddleBoxChainOperations				
/sc/servicechain	List service chains: All or for a specific tenant.	GET	tenant/query deep/query	404
/sc/servicechain	Add a service chain.	POST		404
/sc/servicechain_digest	List name, ID of all service chain tenants.	GET	tenant/query filter/query	404
/sc/servicechain/{id}	Get information about a specified service chain.	GET	tenant/query deep/query id/path	404
/sc/servicechain/{id}	Update a specified service chain.	PUT	id/path	404
/sc/servicechain/{id}	Delete a specified service chain.	DELETE	id/path	404

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
/sc/servicechain/byName/{name}	Get information about the specified service chain.	GET	tenant/query deep/query name/path	404
MiddleBoxChainPolicyOperations				
/mbcp/policy	List waypoint policy: All or for a specified tenant.	GET	tenant/query deep/query	404
/mbcp/policy	Create a waypoint policy.	POST		404
/mbcp/policy_digest	List name, ID of all waypoint policies.	GET	tenant/query	404
/mbcp/policy/{id}	Get information about the specified waypoint policy.	GET	tenant/query deep/query id/path	404
/mbcp/policy/{id}	Update specified waypoint policy.	PUT	id/path	404
/mbcp/policy/{id}	Delete specified waypoint policy.	DELETE	id/path	404
/mbcp/policy/byName/{name}	Get information for the specified waypoint policy.	GET	tenant/query deep/query name/path	404
/mbcp/policy/deploy/{id}	Deploy a waypoint policy.	PUT	id/path	404
/mbcp/policy/undeploy/{id}	Stop using a waypoint policy.	PUT	id/path	404
MiddleBoxOperations				
/mb/middlebox	List all middleboxes.	GET	tenant/query name/path	404
/mb/middlebox	Add a middlebox.	POST		404
/mb/middlebox/{id}	Update a specified middlebox.	PUT	id/path	409
/mb/middlebox/{id}	Delete a specified middlebox.	DELETE	id/path	404
/mb/middlebox/{mbId}	Get information about a specified middlebox.	GET	tenant/query mbId/path	
MiddleBoxTemplateOperations				
/template/import	Import a JSON file containing resource definitions.	POST	uploadedInput-Stream/formdata fileDetail/formdata tenantId/fromdata	

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
/template/wptemplate	Creates a connectivity service template for the specified tenant.	POST		
/template/policy_export/{policyId}	Get information about the specified policy.	GET	tenant/query policyId/path	
Upgrade: /controller/nb/v2/upgrade				
/caches	Displays cache content.	GET		200; 503
/clearCaches	Clears cache content.	GET		200; 503
/inprogress	Returns the status of the upgrade: if it is in progress or is complete.	GET		200; 503
/start	Start the upgrade process.	PUT		200; 503
/status	Returns the status of the upgrade: if it is in progress or is complete.	GET		200; 503
Script Runner: /controller/nb/v2/runscript				
/backup	Downloads the configuration file.	GET		200; 500
/ctlkeystore	Uploads controller private key and controller certificate.	POST	keyInputStream keyFileDetail crtInputStream crtFileDetail	200; 500
/ctltruststore	Downloads the contents of the controller trust store.	GET		200; 500
	Uploads Switch CA Cert / CA Cert to controller trust store.	POST	inputStream fileDetail	200; 500
/doveclientauth	Enables or disables client authentication	POST		200; 500
	Gets the client authentication status.	GET		200; 500
/dovecladd	Uploads CRL (Certification Revocation List) file.	POST		200; 500
/doveclauth	Enables or disables verification against a CRL file.	POST		200; 500
	Gets the current CRL verification status.	GET		200; 500

Table 25. REST API

REST API	Purpose	Operations	Query parameter name/type	Error Code
/doveipsec	Enables or disables IPSec between nodes.	POST		200; 500
	Shows IPSec status.	GET		200; 500
/niststatus	Returns NIST status.	GET		200; 500
	Set NIST status.	POST		200; 500
/ofversion	Shows controller OpenFlow version.	GET		404
	Set OpenFlow version.	PUT		404
/restore	Restores saved configuration.	POST	uploadedInput-Stream fileDetail	200; 500
/runscript	Executes the CLI script with any/variable parameters, and returns the result.	GET		200; 500; 501
/runscript?cmd=run-command%20clear-logs	Clears all logs.			
/ctltruststore/{alias}	Deletes a certificate from the controller truststore.	DELETE		200; 500
/showsecurity/{type}	Displays security files.	GET		200; 500

The following table provides a description of the error codes:

Table 26. Error Code Description

Error Code	Description
200	Operation Successful.
201	Updating Successful.
403	User Authorization Failed.
404	The containerName passed was not found.
405	The FlowGroup does not exist.
406	The Switch Id is null.
408	Referring to non-existent objects.
409	A resource of this type already exists or is in use. The CIDR is already in use.
415	Invalid input data.
500	Failed to delete/modify. No path found between the given two hosts. Host not present/learnt in the network.
503	Service not available.

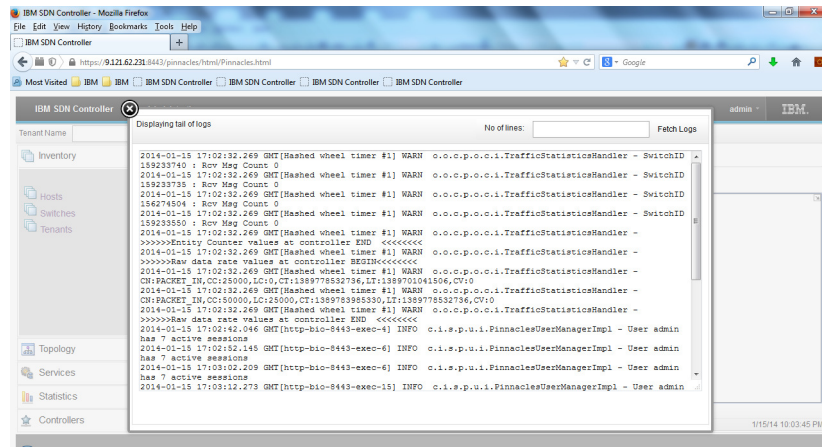
Appendix D. Troubleshooting

Log Information

You can view the complete log files by from the browser interface using the following URL:

`https://<Controller HA external IPv4 address>/log/plog/`

You can also view the tail of the log file by specifying the number of lines.



Network Layers

To help identify a problem, it is important to ensure that the network components are configured and working as expected. There are three network layers in the SDN VE setup:

- Underlay Network
- Management Network
- Overlay Network

Troubleshooting the Underlay Network

The underlay network is configured with Tunnel End Points (TEPs) that connect it with the overlay network.

After configuring the TEP addresses, ensure you can ping the physical network. If you cannot:

- Ensure that the access/trunk port connected to the Host is in the correct VLAN and that tagging is enabled. Only VLAN 1—if used by the DOVE virtual switch—must be untagged.
- Check if the vPort VLAN on the DOVE virtual switch needs to be tagged or untagged.

Troubleshooting the Management Network

Use the SDN VE Controller GUI to ensure that:

- The SDN VE DOVE HA is synchronized.
- The underlay IP subnet is displayed.
- The controller appliance versions should be same on both the controllers.
- The tunnel IP address of the host is displayed under switch information.
- The DCS appliances are correctly listed on the **Service Appliances** page.

Troubleshooting the Overlay Network

A VM assigned to a network profile should be able to see the MAC address of its implicit gateway. If it cannot:

- Check if the networks have been properly exported to the virtual switches. Verify using the `brctl` command on the host.

```
[root@rhel65-27 ~]# brctl show
bridge name      bridge id        STP enabled      interfaces
virbr0           8000.5254005f60d3  yes              virbr0-nic
```

- Check if the networks are configured with the desired IP subnet, mask, and nexthop.
- Check if the DCS config version numbers match. These are displayed under the **Appliances** tab.

Ensure that the VLAN gateway and external gateway configurations are proper.

- Ensure that tagging is correctly configured on both the external switch port connecting to the Physical Host and the virtual switch vPorts.
- Check the VLAN assignment on the edge switch port that connects to the Host.
- Ensure that the gateway appliance NIC is connected to the proper network.
- Ensure VLANs are correctly configured with tagging on the vPort.

Appendix E. Known Issues

The following caveats and limitations were known to exist at the time of initial release for SDN VE version 1.2 and may change as new software becomes available. For the most up-to-date list of known issues, refer to the readme file that is made available with each software update.

IBM SDN VE Controller

- Changing system date or time in the IBM SDN VE Controller demands for restart of the application.
- IBM SDN VE Controller does not support topology where 'multiple hosts with same MAC address' connected from different physical port (ID: XB258906).
- Internal management and control channels are currently not authenticated. (ID: 69071).

Authentication

You must run the system command `restart` on both the primary and secondary controllers every time you change the authentication setting. Restart the secondary controller only after the primary controller GUI comes up after the restart. If both the controllers are restarted simultaneously, the system may become unstable. (ID: XB283669; XB282527)

Broadcast

There is no option to view the Broadcast Tree in the GUI. (ID: XB256535)

CLI

- Tech support dump file takes 5 minutes to save a file with 500 hosts in the Default Network. (ID: XB261661)
- Except the name of the switch, other topology configuration of links etc. cannot be done using CLI.
- System admin has the option to restart the controller. On restart, it asks the admin to re-authenticate. However, it does not check the role after that and shows the wrong view if admin changes to operator.

Cluster Configuration Restore

The restore operation should be performed to clone backed-up configuration onto a fresh installation of the controller. Performing a restore on a controller with existing user configuration will have unpredictable results. (ID: XB292077)

Controller - OpenStack Environment

When using OpenStack, if the you restart the controller, you must reconfigure OpenStack-specific configuration. (ID: XB286904)

External Gateways

- IPv4 addresses in the External Gateway (EGW) pool cannot be modified. Only additions and deletions are supported. (ID: 71464)
- Communication between two tenants is not possible if both tenants map to the same EGW. (ID: 71513)
- EGW fail-over is not triggered when connectivity with the next-hop or 5000V is disrupted.
- When EGW fail-over occurs, existing NAT sessions are not failed restored.
- The last port in the NAT port range will not be used in NAT operations.
- Two tenants that use the same EGW will not be able to communicate with each other. When such a communication is desired, configure two EGWs (one for each tenant) to avoid loopback. (ID: XB290621)
- External Gateway external IP address deletion results in an error. (ID: XB283264)

Resolution: If you need to change the EGW external IP address, you must decommission the existing EGW and deploy a new one.

Flow Management

Setting the priority on the static flow group does not take effect. The priority can be set specifically on each of the contained flows.

Flow Replication & Redirection (FRR)

- Multicast traffic is not supported (ID: XB257205)
- The redirection session shall not get stopped until replica interface goes down (ID: XB254972) and when network admin state goes down (ID: XB254002)

FTP

Problem: FTP from external to overlay via floating IP (or forwarding rule) doesn't work. (ID: XB291028)

In this release, FTP is not supported via floating IP (or forwarding rule).

Gateway Configuration

- **Problem:** Overlay VM cannot ping external VM when a NAT session already exists. (ID: XB290442; XB290431)

Resolution: Current implementation requires a dedicated IP for ping.

GUI

- Host sub-graphs are not getting deleted during auto-refresh.
- Sometimes "Right Click" Action pop ups on topology widget are not displayed on initial right click. The same Right click menus are also available in the Actions menu which can be used.

Layer 3 Service

- L3- Ping between V-Routers is not supported. (ID: XB259546)
- L3 Allow the use of equal cost static routes (ID: XB259476)
- L3 Broadcast is not supported.

Licensing

Localization support is available for displaying the License Agreement. However, in this context, Hungarian language is not supported.

NIST

You must run the system command `restart` on both the primary and secondary controllers every time you change the NIST setting. Restart the secondary controller only after the primary controller GUI comes up after the restart. If both the controllers are restarted simultaneously, the system may become unstable. (ID: XB283669; XB282527)

OpenStack Neutron

Neutron routers with shared networks do not work as expected. In this release, Neutron routers can be used only with dedicated networks. (ID: XB289274)

PKI

You must run the system command `restart` on both the primary and secondary controllers after you upload a certificate or key. Restart the secondary controller only after the primary controller GUI comes up after the restart. If both the controllers are restarted simultaneously, the system may become unstable. (ID: XB283669; XB282527)

Protocols and Traffic

- IGMP reports are not sent to DGW appliances. Manual multicast configuration is needed to overcome this limitation. (ID: 69204)
- Enabling port mirroring results in tagged packets being incorrectly delivered to an un-tagged destination. (ID: 70374)
- Jumbo Frame traffic is not supported by the 5000V switch or DGWs. (ID: 71498)
- Tunnel End-Points (TEPs) cannot be assigned to a user defined VLAN. (ID: 71686)
- FTP server passive mode in networks configured as dedicated cannot be accessed via the EGW. (ID: 71754)

QoS

Problem: 802.1p priority marking from DGW is overwritten when Distributed Gateway (DGW) is connected to SDN VE 5000V Distributed vSwitch port on VLAN 1. (ID: XB283267)

Resolution: Connect the DGW to a switch other than the 5000V Distributed vSwitch. Or, use a VLAN other than 1, if possible.

SDN VE DOVE

- If the HA system remains in failure mode long enough, the primary DMC will revert to “stand-alone” (non-HA) mode to avoid system conflicts. Recovering from this state requires manual intervention to stop HA (system ha stop), re-synchronize (system ha synchronization start), verify (show ha-synchronization), and restart (system ha start).
- DMC HA fails over to two Primaries (ID: XB260869) - This issue occurs when the operating system hosting one of the two nodes of the DMC restarts.

SDN VE HA

- For HA configuration to take effect, the user needs to restart all the controllers in the cluster once (ID: XB259956). This is one time activity, when the controller(s) is / are configured with HA for the very first time.
- The support for OVS with configuration of multiple controllers in HA is not in place. Instead configure single controller.
- The HA configuration needs to be performed in a certain sequence. You must configure the SDN VE HA first. See [“Establish SDN VE Controller HA” on page 43](#). Subsequently, configure the SDN VE DOVE HA. See [“Configure SDN VE DOVE HA” on page 45](#). (ID:XB276068)

Topology

Topology links timeout if the switches do not respond to LLDP discovery packets within a time limit of 30 seconds (ID: XB223202)

Virtual Machines (VMs)

- VM that do not participate in network traffic may not appear in the show endpoints output. (ID: 69562)
- If VMs are not restarted after the host server is power cycled or rebooted, entries for those VMs will still appear in the show endpoints output. (ID: 71538)

Virtual Switching

- Interface level configuration of ports in the 5000V vDS are not supported, though not explicitly disallowed in the CLI. (ID: 71003)
- Only one 5000V vDS per host server is supported for SDN VE use. (ID: 71750)
- **Problem:** 5000V ports are blocked due to address registration errors. (ID: XB283427)
Resolution: Any of the following steps may help to resolve the issue:
 - Disconnect and reconnect the port
 - Disconnect and reconnect the TEP IP vmkernel interface
 - Disconnect and reconnect the ESX host from the 5000V Distributed vSwitch

VLAN Gateways

- VLAN gateway packet forwarding disrupted when Tunnel-IP's next hop connectivity fails (ARP resolve fail)

VM IP Addresses

- If an IPv4 address is removed from a VM, it is not unregistered if there is continuous traffic for it. (ID: 71502)
- IPv4 address conflict on the same host and same virtual network results in loss of communication. (ID: 71560)
- If a VM's IPv4 address is changed, the IPv4 address is not unregistered if there is continued traffic to another VM on same host. (ID: 71604)
- IPv6 addresses are not presently supported.

VRRP

- The priority value configured on a DGW is not effective. DGW with the smaller IP address is always selected as the Master. (ID: XB290858)
- DGW VRRP failover may take more than 10 seconds. (ID: XB283557)

Waypoint Connectivity Service

High-Availability

Failover of a routed Waypoint device from active to a standby may result in traffic disruption of more than two minutes. (ID: XB281976)

Workaround: Configure the Waypoint device in the transparent mode with the SDN VE implicit gateway providing the routing functionality instead of the Waypoint device.

Routed NAT Devices

Waypoint devices configured in routed NAT mode replace the original IP of an incoming packet. These devices should not be shared between service chains because endpoints from one service chain may communicate with endpoints from another service because of the IP address replacement. (ID: XB281991; 14422)

For example:

If a routed NAT Waypoint (Wnat) is used in two service chains (S3 and S4) as follows:

C1 → Wnat → S3 and C2 → Wnat → S4

Endpoints from C1 may be able to reach endpoints in S4, and endpoints in C2 may reach endpoints in S3.

To avoid this, configure two routed NAT Waypoints—one for each service chain—to ensure traffic is properly segregated.

Service Chains

- Service chains cannot be configured between dedicated and shared groups. (ID: XB283248)
- Service chains between shared groups can only be configured from the administrator tenant i.e. DOVE admin. (ID: XB283248)

Appendix F. Upgrading IBM SDN VE Components

IBM SDN VE Controller

The IBM SDN VE controller can be upgraded using the CLI or GUI.

Note: The upgrade image should be placed on a web server that is accessible to the controller. The image file extension will be `.img`.

CLI

Use the `system upgrade` command. See [“system upgrade” on page 254](#) for the command details.

GUI

1. Login to the controller GUI using `https://<Controller IP address>:8443`.
2. Select **Administration > Upgrade**. The **Upgrade** window is displayed.
3. Specify the location of the image file in the **Image Path** field.
4. Select **Start Upgrade**.

The upgrade status can be viewed in the **Upgrade Status** section of the window.

DOVE Connectivity Service (DCS)

If an upgrade image (.img) file is available:

1. Reset the role of that DCS node.
2. Upgrade the DSA and set the role once the DSA boots up.

This will have to be done one DSA node at a time and there must be at least two DCS nodes to prevent loss of connectivity.

If there is no upgrade image available:

1. Reset the role of a DCS node
2. Remove it from the network
3. Deploy a new DSA node (with the new image)
4. Add it to the DMC and assign the role.

This again will have to be done one DSA node at a time.

DOVE Gateway (DGW)

To avoid service outage during upgrade there should be at least two DGW assigned per VNID.

If an upgrade image (.img) file is available:

1. First reset the role of DGW node.
2. Upgrade and set the role back when DSA boots.

If there is no upgrade image available:

1. Reset DGW role of retiring node (this step will force user to manually delete that DGW configuration).
2. Deploy new DSA node and set role.
3. Configure new DGW node.

Appendix G. Getting Help and Technical Assistance

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your system, and whom to call for service, if it is necessary.

Before You Call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system and any optional devices are turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system. Information about diagnostic tools is in the *Problem Determination and Service Guide* on the IBM Documentation CD that comes with your system.
- Go to the IBM support website at <http://www.ibm.com/systems/support/> to check for technical information, hints, tips, and new device drivers or to submit a request for information.

You can solve many problems without outside assistance by following the troubleshooting procedures that IBM provides in the online help or in the documentation that is provided with your IBM product. The documentation that comes with IBM systems also describes the diagnostic tests that you can perform. Most systems, operating systems, and programs come with documentation that contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the documentation for the operating system or program.

Using the Documentation

Information about your IBM system and pre-installed software, if any, or optional device is available in the documentation that comes with the product. That documentation can include printed documents, online documents, ReadMe files, and Help files. See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software. IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to <http://www.ibm.com/systems/support/> and follow the instructions. Also, some documents are available through the IBM Publications Center at <http://www.ibm.com/shop/publications/order/>.

Getting Help and Information on the World Wide Web

On the World Wide Web, the IBM website has up-to-date information about IBM systems, optional devices, services, and support. The address for IBM Systems information is <http://www.ibm.com/systems>.

You can find service information for IBM systems and optional devices at <http://www.ibm.com/support/>.

Software Service and Support

Through IBM Support Line, you can get telephone assistance, for a fee, with usage, configuration, and software problems with SDN VE. For information about which products are supported by Support Line in your country or region, see <http://www.ibm.com/services/sl/products/>.

For more information about Support Line and other IBM services, see <http://www.ibm.com/services/>, or see <http://www.ibm.com/planetwide/> for support telephone numbers. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

Hardware Service and Support

You can receive hardware service through your IBM reseller or IBM Services. To locate a reseller authorized by IBM to provide warranty service, go to <http://www.ibm.com/partnerworld/> and click **Find Business Partners** on the right side of the page. For IBM support telephone numbers, see <http://www.ibm.com/planetwide/>. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

IBM Taiwan Product Service

台灣 IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路 7 號 3 樓
電話：0800-016-888

IBM Taiwan product service contact information:

IBM Taiwan Corporation
3F, No 7, Song Ren Rd.
Taipei, Taiwan
Telephone: 0800-016-888

Appendix H. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM 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:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

INTERNATIONAL BUSINESS MACHINES CORPORATION 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 states 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. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM 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 IBM product, and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Trademarks

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. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc., in the United States, other countries, or both and is used under license therefrom.

Intel, Intel Xeon, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc., in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Important Notes

Processor speed indicates the internal clock speed of the microprocessor; other factors also affect application performance.

CD or DVD drive speed is the variable read rate. Actual speeds vary and are often less than the possible maximum.

When referring to processor storage, real and virtual storage, or channel volume, KB stands for 1024 bytes, MB stands for 1 048 576 bytes, and GB stands for 1 073 741 824 bytes.

When referring to hard disk drive capacity or communications volume, MB stands for 1 000 000 bytes, and GB stands for 1 000 000 000 bytes. Total user-accessible capacity can vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard disk drive bays with the largest currently supported drives that are available from IBM.

Maximum memory might require replacement of the standard memory with an optional memory module.

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software might differ from its retail version (if available) and might not include user manuals or all program functionality.

Particulate Contamination

Attention: Airborne particulates (including metal flakes or particles) and reactive gases acting alone or in combination with other environmental factors such as humidity or temperature might pose a risk to the devices that run the software described in this document. Risks that are posed by the presence of excessive particulate levels or concentrations of harmful gases include damage that might cause the device to malfunction or cease functioning altogether. This specification sets forth limits for particulates and gases that are intended to avoid such damage. The limits must not be viewed or used as definitive limits, because numerous other factors, such as temperature or moisture content of the air, can influence the impact of particulates or environmental corrosives and gaseous contaminant transfer. In the absence of specific limits that are set forth in this document, you must implement practices that maintain particulate and gas levels that are consistent with the protection of human health and safety. If IBM determines that the levels of particulates or gases in your environment have caused damage to the device, IBM may condition provision of repair or replacement of devices or parts on implementation of appropriate remedial measures to mitigate such environmental contamination. Implementation of such remedial measures is a customer responsibility.

Contaminant	Limits
Particulate	<ul style="list-style-type: none">• The room air must be continuously filtered with 40% atmospheric dust spot efficiency (MERV 9) according to ASHRAE Standard 52.2¹.• Air that enters a data center must be filtered to 99.97% efficiency or greater, using high-efficiency particulate air (HEPA) filters that meet MIL-STD-282.• The deliquescent relative humidity of the particulate contamination must be more than 60%².• The room must be free of conductive contamination such as zinc whiskers.
Gaseous	<ul style="list-style-type: none">• Copper: Class G1 as per ANSI/ISA 71.04-1985³• Silver: Corrosion rate of less than 300 Å in 30 days
<p>¹ ASHRAE 52.2-2008 - <i>Method of Testing General Ventilation Air-Cleaning Devices for Removal Efficiency by Particle Size</i>. Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.</p> <p>² The deliquescent relative humidity of particulate contamination is the relative humidity at which the dust absorbs enough water to become wet and promote ionic conduction.</p> <p>³ ANSI/ISA-71.04-1985. <i>Environmental conditions for process measurement and control systems: Airborne contaminants</i>. Instrument Society of America, Research Triangle Park, North Carolina, U.S.A.</p>	

Documentation Format

The publications for this product are in Adobe Portable Document Format (PDF) and should be compliant with accessibility standards. If you experience difficulties when you use the PDF files and want to request a web-based format or accessible PDF document for a publication, direct your mail to the following address:

Information Development
IBM Corporation
205/A015
3039 E. Cornwallis Road
P.O. Box 12195
Research Triangle Park, North Carolina 27709-2195
U.S.A.

In the request, be sure to include the publication title.

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.