Firefly Perimeter Getting Started Guide for VMware

Published: 2014-02-16

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Firefly Perimeter Getting Started Guide for VMware

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at http://www.juniper.net/support/eula.html. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at http://www.juniper.net/techpubs/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at http://www.juniper.net/books.

Documentation Conventions

Table 1 on page x defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.

Table 2 on page x defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric** *metric***>;**

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [*community-ids*]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract,

or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://www.juniper.net/alerts/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://tools.juniper.net/SerialNumberEntitlementSearch/

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see http://www.juniper.net/support/requesting-support.html.

PART 1

Overview

CHAPTER 1

Firefly Perimeter Overview

Understanding Firefly Perimeter

Firefly Perimeter is a virtual security appliance that provides security and networking services at the perimeter or edge in virtualized private or public cloud environments. Firefly Perimeter runs as a virtual machine (VM) on a standard x86 server.

Firefly Perimeter enables advanced security and routing at the network edge in a multitenant virtualized environment. Firefly Perimeter is built on Junos OS and delivers similar networking and security features available on SRX Series devices for the branch.

Some of the key benefits of Firefly Perimeter in virtualized private or public cloud multitenant environments include:

- Stateful firewall protection at the tenant edge
- Faster deployment of virtual firewalls
- Full routing, Virtual Private Network (VPN) and networking capabilities
- Complementary with the Juniper Networks Firefly Host for inter-VM security
- Centralized and local management

Related Documentation

- Specifications for Firefly Perimeter Installation on page 27
- Firefly Perimeter Basic Settings on page 28
- Installation Requirements for Firefly Perimeter with VMware on page 29

Features Supported on Firefly Perimeter with VMware

Firefly Perimeter inherits many features from the SRX Series product line. However, because some SRX Series features are not directly applicable in a virtualized environment, they have been excluded from the Firefly Perimeter product line. Table 3 on page 4 describes the available features on Firefly Perimeter as of Junos OS Release 12.1X46-D10. For feature roadmap details, contact your Juniper Networks representative.

Table 3: Features Supported on Firefly Perimeter

Feature	Support on Firefly Perimeter
Address Books and Address Sets:	
Address books	Yes
Address sets	Yes
Global address objects or sets	Yes
Nested address groups	Yes
Administrator Authentication:	
Local authentication	Yes
RADIUS	Yes
TACACS+	Yes
Alarms:	
Chassis alarms	Yes
Interface alarms	Yes
System alarms	Yes
Application Layer Gateways:	
DNS ALG	Yes
DNS doctoring support	No
DNS, FTP, RTSP, and TFTP ALGs (Layer 2) with chassis clustering	No
DSCP marking for SIP, H.323, MGCP, and SCCP ALGs	Yes
FTP	Yes
H.323	Yes
Avaya H.323	Yes
IKE	Yes
MGCP	Yes
PPTP	Yes
RSH	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
RTSP	Yes
SCCP	Yes
SIP	Yes
SIP ALG–NEC	Yes
SQL	Yes
MS RPC	Yes
SUN RPC	Yes
TALK	Yes
TFTP	Yes
Attack Detection and Prevention:	
Bad IP option	Yes
Block fragment traffic	Yes
FIN flag without ACK flag set protection	Yes
ICMP flood protection	Yes
ICMP fragment protection	Yes
IP address spoof	Yes
IP address sweep	Yes
IP record route option	Yes
IP security option	Yes
IP stream option	Yes
IP strict source route option	Yes
IP timestamp option	Yes
Land attack protection	Yes
Large size ICMP packet protection	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
Loose source route option	Yes
Ping of death attack protection	Yes
Port scan	Yes
Source IP-based session limit	Yes
SYN-ACK-ACK proxy protection	Yes
SYN and FIN flags set protection	Yes
SYN flood protection	Yes
SYN fragment protection	Yes
TCP address sweep	Yes
TCP packet without flag set protection	Yes
Teardrop attack protection	Yes
UDP address sweep	Yes
UDP flood protection	Yes
Unknown IP protocol protection	Yes
Whitelist for SYN flood screens	Yes
WinNuke attack protection	Yes
Autoinstallation:	
Autoinstallation	Yes
Class of Service:	
Classifiers	Yes
Code-point aliases	Yes
Egress interface shaping	Yes
Forwarding classes	Yes
High-priority queue on Services Processing Card	No

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
Ingress interface policer	Yes
Schedulers	Yes
Simple filters	Yes
Transmission queues	Yes
Tunnels NOTE: GRE and IP-IP tunnels only.	Yes
Virtual channels	Yes
Diagnostics Tools:	
CLI terminal	Yes
Flow monitoring cflowd version 5 and flow monitoring cflowd version 8	Yes
Flow monitoring cflowd version 9	No
Ping host	Yes
Ping MPLS	Yes
Traceroute	Yes
Ping Ethernet (CFM)	No
Traceroute Ethernet (CFM)	No
DNS Proxy:	
DNS proxy cache	Yes
DNS proxy with split DNS	Yes
Dynamic DNS	No
Dynamic Host Configuration Protocol:	
DHCPv6 client	No
DHCPv4 client	Yes
DHCPv6 relay agent	No

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
DHCPv4 relay agent	Yes
DHCPv6 server	Yes
DHCPv4 server	Yes
DHCP server address pools	Yes
DHCP server static mapping	Yes
Ethernet Link Aggregation:	
Routing mode:	
LACP in chassis cluster pair	No
LACP in standalone device	No
Layer 3 LAG on routed ports	No
Static LAG in chassis cluster mode	No
Static LAG in standalone mode	No
Ethernet Link Fault Management:	
Interfaces supported:	
LACP in chassis cluster pair	No
LACP in standalone mode	No
Static LAG in chassis cluster mode	No
Static LAG in standalone mode	No
Physical interface (encapsulations):	
ethernet-ccc	No
extended-vlan-ccc	No
ethernet-tcc	No
extended-vlan-tcc	No
Interface family:	

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
inet	Yes
mpls	Yes
ccc	No
tcc	No
iso	Yes
ethernet-switching	No
inet6	Yes
Aggregated Ethernet interface:	
Static LAG	No
LACP enabled LAG	No
Interface family:	
ethernet-switching	No
inet	Yes
inet6	Yes
iso	Yes
mpls	Yes
File Management:	
Clean up unnecessary files	Yes
Delete backup software image	Yes
Delete individual files	Yes
Download system files	Yes
Encrypt/decrypt configuration files	Yes
Manage account files	Yes
Rescue	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
System zeroize	Yes
Monitor start	Yes
Archive files	Yes
Calculate checksum	Yes
Compare files	Yes
Rename files	Yes
Firewall Authentication:	
Firewall authentication on Layer 2 transparent authentication	No
LDAP authentication server	Yes
Local authentication server	Yes
Pass-through authentication	Yes
RADIUS authentication server	Yes
SecurID authentication server	Yes
Web authentication	Yes
Flow-Based and Packet-Based Processing:	
Alarms and auditing	Yes
End-to-end packet debugging	No
Flow-based processing	Yes
Network processor bundling	No
Packet-based processing	Yes
Selective stateless packet-based services	Yes
Interfaces:	
Physical and Virtual Interface:	
Ethernet interface	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
Gigabit Ethernet interface	Yes
Services:	
Aggregated Ethernet interface	No
GRE interface	Yes
IEEE 802.1X dynamic VLAN assignment	No
IEEE 802.1X MAC bypass	No
IEEE 802.1X port-based authentication control with multisupplicant support	No
Interleaving using MLFR	No
Internally configured interface used by the system as a control path between the WXC Integrated Services Module and the Routing Engine	No
Internally generated GRE interface (gr-0/0/0)	Yes
Internally generated IP-over-IP interface (ip-0/0/0)	Yes
Internally generated link services interface	Yes
Internally generated Protocol Independent Multicast de-encapsulation interface	Yes
Internally generated Protocol Independent Multicast encapsulation interface	Yes
Link fragmentation and interleaving interface	Yes
Link services interface	Yes
Loopback interface	Yes
Management interface	Yes
PPP interface	No
PPPoE-based radio-to-router protocol	No
PPPoE interface	No

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
Promiscuous mode on interfaces	Yes NOTE: Promiscuous mode needs to be enabled on hypervisor.
Secure tunnel interface	Yes
IP Monitoring:	
IP monitoring with route failover (for standalone devices and redundant Ethernet interfaces)	Yes
IP monitoring with interface failover (for standalone devices)	Yes
Track IP enhancements (IP Monitoring using RPM)	No
IP Security:	
Acadia - Clientless VPN	No
Alarms and auditing	Yes
Antireplay (packet replay attack prevention)	Yes
Authentication	Yes
Authentication Header (AH)	Yes
Autokey management	Yes
Automated certificate enrollment using SCEP	Yes
Automatic generation of self-signed certificates	Yes
Bridge domain and transparent mode	No
Certificate - Configure local certificate sent to peer	Yes
Certificate - Configure requested CA of peer certificate	Yes
Certificate - Encoding: PKCS7, X509, PEM, DERs	Yes
Certificate - RSA signature	Yes
Chassis clusters (active/backup and active/active)	Yes
Class of service	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
CRL update at user-specified interval	Yes
Config Mode (draft-dukes-ike-mode-cfg-03)	Yes
Dead peer detection (DPD)	Yes
Diffie-Hellman (PFS) Group 1	Yes
Diffie-Hellman (PFS) Group 2	Yes
Diffie-Hellman (PFS) Group 5	Yes
Diffie-Hellman Group 1	Yes
Diffie-Hellman Group 2	Yes
Diffie-Hellman Group 5	Yes
Digital signature generation	Yes
Dynamic IP address	Yes
Dynamic IPsec VPNs	No
Encapsulating Security Payload (ESP) protocol	Yes
Encryption algorithms 3DES	Yes
Encryption algorithms AES 128, 192, and 256	Yes
Encryption algorithms DES	Yes
Encryption algorithms NULL (authentication only)	Yes
Entrust, Microsoft, and Verisign certificate authorities (CAs)	Yes
External Extended Authentication (Xauth) to a RADIUS server for remote access connections	Yes
Group Encrypted Transport (GET VPN)	No
Group VPN with dynamic policies	No
Hard lifetime limit	Yes
Hardware IPsec (bulk crypto) Cavium/RMI	No

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
Hash algorithms MD5	Yes
Hash algorithms SHA-1	Yes
Hash algorithms SHA-2 (SHA-256)	Yes
Hub & spoke VPN	Yes
Idle timers for IKE	Yes
Improvements in VPN debug capabilities	Yes
Initial contact	Yes
Invalid SPI response	Yes
IKE Diffie-Hellman Group 14 support	Yes
IKE Phase 1	Yes
IKE Phase 1 lifetime	Yes
IKE Phase 2	Yes
IKE Phase 2 lifetime	Yes
IKE and IPsec predefine proposal sets to work with dynamic VPN client	No
IPsec tunnel termination in routing-instances	Yes
IKE support	Yes
IKEv1	Yes
IKEv1 authentication, preshared key	Yes
IKEv2	Yes
Local IP address management - VPN XAuth support	Yes
Local IP address management support for DVPN	No
Manual installation of DER-encoded and PEM-encoded CRLs	Yes
Manual key management	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
Manual proxy-ID (Phase 2 ID) configuration	Yes
NHTB - Next Hop Tunnel Binding	Yes
New IPsec Phase 2 authentication algorithm	Yes
Online CRL retrieval through LDAP and HTTP	Yes
Package dynamic VPN client	No
Policy-based VPN	Yes
Preshared key (PSK)	Yes
Prioritization of IKE packet processing	Yes
Reconnect to dead IKE peer	Yes
Remote access	Yes
Remote access user IKE peer	Yes
Remote access user-group IKE peer - group IKE ID	Yes
Route-based VPN	Yes
SHA-2 IPsec support	Yes
Soft lifetime	Yes
Static IP address	Yes
Suites: standard, compatible, basic, and custom-created	Yes
Support for NHTB when the st0.x interface is bound to a routing instance	Yes
Support for remote access peers with shared IKE identity + mandatory XAuth	Yes
Support group IKE IDs for dynamic VPN configuration	No
TOS/DSCP honoring/coloring (inner/outer)	Yes
Tunnel mode with clear/copy/set Don't Fragment bit	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
UAC Layer 3 enforcement	Yes
Virtual router support for route-based VPNs	Yes
VPN monitoring (proprietary)	Yes
X.509 encoding for IKE	Yes
XAuth (draft-beaulieu-ike-xauth-03)	Yes
IPv6 Support:	
Flow-based forwarding and security features:	
Advanced flow	Yes
DS-Lite concentrator (aka AFTR)	No
DS-Lite initiator (aka B4)	No
Firewall filters	Yes
Forwarding option: flow mode	Yes
Multicast flow	Yes
Screens	Yes
Security policy (firewall)	Yes
Security policy (IDP)	No
Security policy (user role firewall)	No
Zones	Yes
IPv6 ALG support for FTP: Routing, NAT, NAT-PT support	Yes
IPv6 ALG support for ICMP: Routing, NAT, NAT-PT support	Yes
IPv6 NAT: NAT-PT, NAT support	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
IPv6 NAT64	Yes
IPv6–related protocols: BFD, BGP, ECMPv6, ICMPv6, ND, OSPFv3, RIPng	Yes
IPv6 ALG support for TFTP	Yes
System services: DHCPv6, DNS, FTP, HTTP, ping, SNMP, SSH, syslog, Telnet, traceroute	Yes
Packet-based forwarding and security features:	
Class of service	Yes
Firewall filters	Yes
Forwarding option: packet mode	Yes
Chassis cluster	
Active-active	Yes
Active-passive	Yes
Multicast flow	Yes
IPv6 IP Security:	
4in4 and 6in6 policy-based site-to-site VPN, AutoKey IKEv1	No
4in4 and 6in6 policy-based site-to-site VPN, manual key	No
4in4 and 6in6 route-based site-to-site VPN, AutoKey IKEv1	No
4in4 and 6in6 route-based site-to-site VPN, manual key	No
Log File Formats:	
System (control plane) log file formats:	
Binary format (binary)	No
Structured syslog (sd-syslog)	Yes
Syslog (syslog)	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
WebTrends Enhanced Log Format (WELF)	No
Security (data plane) log file formats:	
Binary format (binary)	Yes
Structured syslog (sd-syslog)	Yes
Syslog (syslog)	Yes
WebTrends enhanced log format (WELF)	Yes
MPLS:	
CCC and TCC	No
CLNS	Yes
Interprovider and carrier-of-carriers VPNs	Yes
Layer 2 VPNs for Ethernet connections	Yes NOTE: Promiscuous mode needs to be enabled on hypervisor.
Layer 3 MPLS VPNs	Yes
LDP	Yes
MPLS VPNs with VRF tables on provider edge routers	Yes
Multicast VPNs	Yes
OSPF and IS-IS traffic engineering extensions	Yes
P2MP LSPs	Yes
RSVP	Yes
Secondary and standby LSPs	Yes
Standards-based fast reroute	Yes
Multicast:	
Filtering PIM register messages	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
IGMP	Yes
PIM RPF routing table	Yes
Primary routing mode (dense mode for LAN and sparse mode for WAN)	Yes
Protocol Independent Multicast Static RP	Yes
Session Announcement Protocol (SAP)	Yes
SDP	Yes
Multicast VPN:	
Basic multicast features in C-instance	Yes
Multicast VPN membership discovery with BGP	Yes
P2MP LSP support	Yes
P2MP OAM - P2MP LSP ping	Yes
Reliable multicast VPN routing information exchange	Yes
Network Address Translation:	
Destination IP address translation	Yes
Disabling source NAT port randomization	Yes
Interface source NAT pool port	Yes
NAT address pool utilization threshold status	Yes
NAT traversal (NAT-T) for site-to-site IPsec VPNs (IPv4)	Yes
Persistent NAT	Yes
Persistent NAT binding for wildcard ports	Yes
Persistent NAT hairpinning	Yes
Maximize persistent NAT bindings	No
Pool translation	Yes
Proxy ARP (IPv4)	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
Proxy NDP (IPv6)	Yes
Removing persistent NAT query bindings	Yes
Rule-based NAT	Yes
Rule translation	Yes
Source address and group address translation for multicast flows	Yes
Source IP address translation	Yes
Static NAT	Yes
Network Operations and Troubleshooting:	
Event policies	Yes
Event scripts	Yes
Operation scripts	Yes
XSLT commit scripts	Yes
Network Time Protocol:	
NTP support	Yes
Packet Capture:	
Packet capture NOTE: Packet capture, in this context, refers to standard interface packet capture. It is not part of the IDP. Packet capture is supported only on physical interfaces and tunnel interfaces; for example, *gr*, *ip*, *st0*, *lsq-/ls-*. Packet capture is not supported on redundant Ethernet interfaces (*reth*).	Yes
Real-Time Performance Monitoring Probe:	
RPM probe	Yes
One-way timestamps	Yes
Routing:	
BGP	Yes
BGP extensions for IPv6	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
BGP Flowspec	No
Compressed Real-Time Transport Protocol (CRTP)	No
ECMP flow-based forwarding	No
Internet Group Management Protocol (IGMP)	Yes
IPv4 options and broadcast Internet diagrams	Yes
IPv6 routing, forwarding, global address configuration, and Internet Control Message Protocol (ICMP)	Yes
IS-IS	Yes
Multiple virtual routers	Yes
Neighbor Discovery Protocol (NDP) and Secure NDP	Yes
OSPF v2	Yes
OSPF v3	Yes
RIP next generation (RIPng)	Yes
RIP v1, v2	Yes
Static routing	Yes
Virtual Router Redundancy Protocol (VRRP)	Yes
Secure Web Access:	
CAs	Yes
HTTP	Yes
HTTPS	Yes
Security Policy Support:	
Address books/address sets	Yes
Custom policy applications	Yes
Global policy	Yes
Policy application timeouts	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
Policy applications and application sets	Yes
Policy hit-count tracking	Yes
Schedulers	Yes
Security policies for self-traffic	Yes
SSL proxy	No
User role firewall	No
Common predefined applications	Yes
Shadow policy	Yes
Security Zone:	
Functional zone	Yes
Security zone	Yes
Session Logging:	
Accelerating security and traffic logging	Yes
Aggressive session aging	Yes
Getting information about sessions	Yes
Logging to a single server	Yes
Session logging with NAT information	Yes
SMTP:	
SMTP support	Yes
SNMP:	
SNMP support	Yes
Stateless Firewall Filters:	
Stateless firewall filters (ACLs)	Yes
Stateless firewall filters (simple filter)	No
System Log Files:	

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
Archiving system logs	Yes
Configuring system log messages	Yes
Disabling system logs	Yes
Filtering system log messages	Yes
Multiple system log servers (control-plane logs)	Yes
Sending system log messages to a file	Yes
Sending system log messages to a user terminal	Yes
Viewing data plane logs	Yes
Viewing system log messages	Yes
Upgrading and Rebooting:	
Autorecovery	No
Boot device configuration	No (N.A.)
Boot device recovery	No (N.A.)
Chassis components control	Yes
Chassis restart	Yes
Download manager	Yes
Dual-root partitioning	No
In-band cluster upgrade	No
Low-impact cluster upgrades	No
Software upgrades and downgrades	Yes
User Interfaces:	
CLI	Yes
J-Web user interface	Yes
Junos XML protocol	Yes

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
Network and Security Manager	No
Junos Space Security Director	Yes
SRC application	No
Junos Space Virtual Director	Yes
Authentication with IC Series Devices	
Captive Portal	Yes
Junos OS Enforces in UAC deployments	Yes
Chassis Cluster Support on VMWare	
Active/active chassis cluster	Yes
ALGs	Yes
Chassis cluster formation	Yes
Control plane failover	Yes
Dampening time between back-to-back redundancy group failover	Yes
Data plane failover	Yes
Dual control links	No
Dual fabric links	Yes
In-band cluster upgrade	No
Junos OS flow-based routing functionality	Yes
Layer 2 Ethernet switching capacity	No
Layer 2 LAG	No
Layer 3 LAG	No
LACP support for Layer 2	No
LACP support for Layer 3	No
Low-impact cluster upgrade (ISSU Light)	No

Table 3: Features Supported on Firefly Perimeter *(continued)*

Feature	Support on Firefly Perimeter
Low latency firewall	No
Multicast routing	Yes
PPPoE over redundant Ethernet interface	No
Redundant Ethernet interfaces	Yes
Redundant Ethernet interface LAGs	No
Redundant Ethernet or aggregate Ethernet interface monitoring	Yes
Redundancy group 0 (backup for Routing Engine)	Yes
Redundancy group 1 through 128	Yes
Upstream device IP address monitoring	Yes
Upstream device IP address monitoring on a backup interface	Yes
Chassis Management Support	
Chassis management	Yes
VPLS	
Filtering and Policing (Packet-Based)	Yes

Table 4 on page 25 lists additional features that are not supported on Firefly.

Table 4: Firefly Feature Support Information

Feature	Firefly
Application Identification (Junos OS)	No
Authentication with IC Series Devices	No
General Packet Radio Service	No
Intrusion Detection and Prevention	No
Layer 2 Mode	No
Logical Systems	No
Power over Ethernet	No

Table 4: Firefly Feature Support Information *(continued)*

Feature	Firefly
Public Key Infrastructure	No
Remote Device Access	No
Route Reflector	No
RPM Probe	No
Services Offloading	No
Transparent Mode	No
Unified Threat Management	No
USB Modem	No
Voice over Internet Protocol with Avaya	No
Wireless Local Area Network	No
Group VPN	No
Multicast for AutoVPN	No
Dynamic VPN (DVPN).	No

Related Documentation

- Understanding Firefly Perimeter on page 3
- Specifications for Firefly Perimeter Installation on page 27
- Firefly Perimeter Basic Settings on page 28

CHAPTER 2

System Requirements

Specifications for Firefly Perimeter Installation

Table 5 on page 27 lists the specifications for Firefly Perimeter.

Table 5: Specifications for Firefly Perimeter

Component	Specification
Memory	2 GB
Disk space	2 GB
vCPUs	2
vNICs	Up to 10
Virtual Network Interface Card type (NIC)	E1000

Table 6 on page 27 lists the hardware specifications for the host machine that runs Firefly Perimeter VM.

Table 6: Hardware Specifications for Host Machine

Component	Specification
Host memory size	Minimum 4 GB
Host processor type	x86_64

NOTE:

- Ensure that the physical server includes multi-core CPU.
- The Host machine must support VMware.

  For the Hardware Compatibility List, see:

  http://vmware.com.

Related Documentation

- Understanding Firefly Perimeter on page 3
- Firefly Perimeter Basic Settings on page 28
- Installation Requirements for Firefly Perimeter with VMware on page 29

Firefly Perimeter Basic Settings

Firefly Perimeter is a security device that requires these basic configuration settings to function:

- Interfaces must be assigned IP addresses.
- Interfaces must be bound to zones.
- Policies must be configured between zones to permit or deny traffic.

Firefly Perimeter has the following default configurations set when you power it on for the first time.

Table 7 on page 28 lists the basic settings for interfaces.

Table 7: Basic Settings for Interfaces

Interface	Security Zones	DHCP State
ge-0/0/0	trust	client
ge-0/0/1 to ge-0/0/3	trust	server

Table 8 on page 28 lists the basic settings for the security policies.

Table 8: Basic Settings for Security Policies

Source Zone	Destination Zone	Policy Action
trust	untrust	permit
trust	trust	permit
untrust	trust	deny

Table 9 on page 29 lists the basic settings for the NAT rule.

Table 9: Basic Settings for NAT Rule

Source Zone	Destination Zone	Policy Action
trust	untrust	source NAT to untrust zone interface

Related Documentation

- Understanding Firefly Perimeter on page 3
- Specifications for Firefly Perimeter Installation on page 27
- Installation Requirements for Firefly Perimeter with VMware on page 29

Installation Requirements for Firefly Perimeter with VMware

Table 10 on page 29 lists the supported version of VMware Hypervisor.

Table 10: Supported Version of VMware hypervisor

VMware Hypervisor	Hypervisor Version
VMware vSphere ESXi	5.0 and 5.1

NOTE: Create an account on the VMware website at http://vmware.com to access the downloads and to obtain the license key for VMware.

Related Documentation

- Understanding Firefly Perimeter on page 3
- Specifications for Firefly Perimeter Installation on page 27
- Firefly Perimeter Basic Settings on page 28

PART 2

Installation

CHAPTER 3

Firefly Perimeter Installation and Connection

Installing Firefly Perimeter with VMware vSphere Client

NOTE: **The following installation steps were performed by connecting VMware vSphere Client 5.0 directly to a host.**

To install Firefly Perimeter with VMware vSphere Client:

1. Download and install the VMware vSphere Client available at:

   http://vmware.com.

2. On the same computer, download the Firefly Perimeter software package available at:

   http://www.juniper.net/support/downloads/.

3. Start the VMware vSphere Client and log in with your credentials. See Figure 1 on page 34.

Figure 1: VMware vSphere Client Login

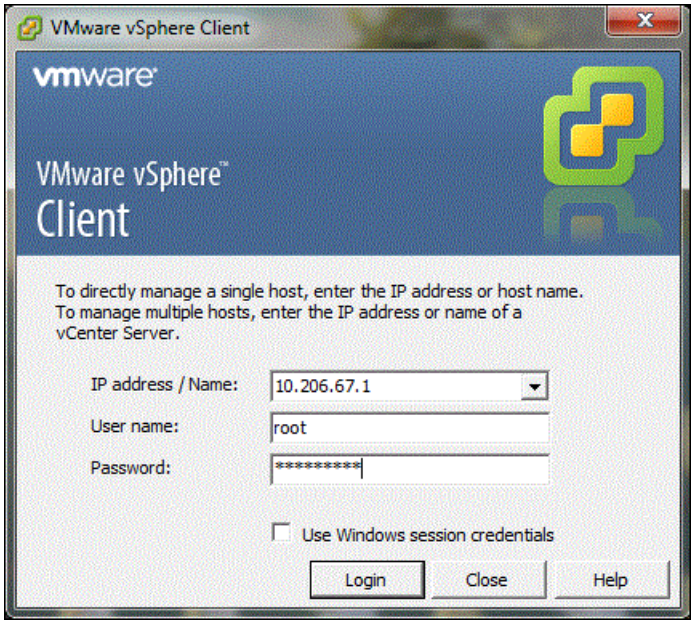

4. Click **File** > **Deploy OVF Template** as shown in Figure 2 on page 34.

Figure 2: OVF Template

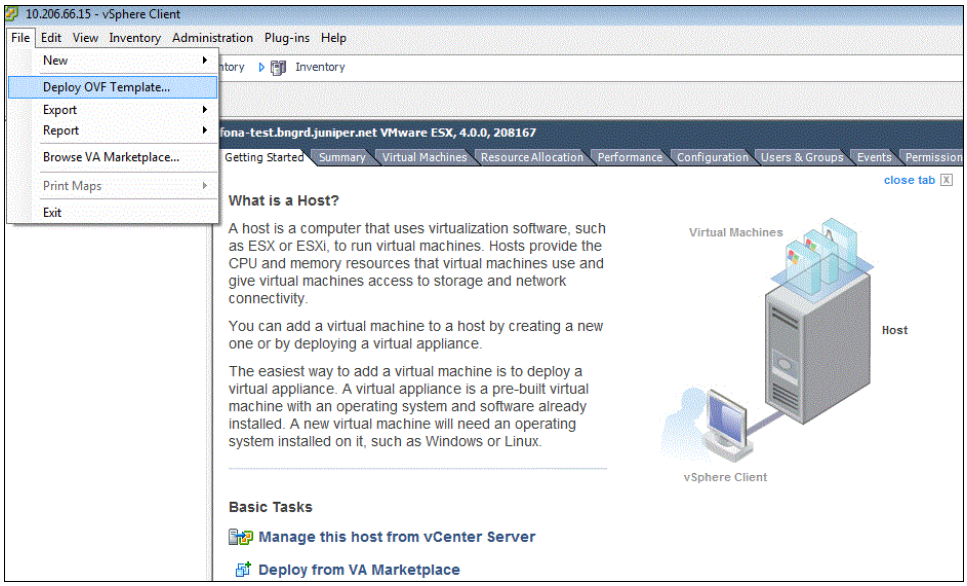

5. Click **Browse** to locate the Firefly Perimeter software package and then click **Next**.
6. Click **Next** on the OVF Template Details window, to proceed with the installation. Click **Cancel** to discard the most recent change.
7. Click **Accept** > **Next** on the End User License Agreement window, to proceed with the installation.

8. Change the default Firefly Perimeter VM name appropriately in the Name box and click **Next**. See Figure 3 on page 35. It is advisable to keep this name the same as the hostname you intend to give to your VM.

Figure 3: Changing the Firefly Perimeter Name

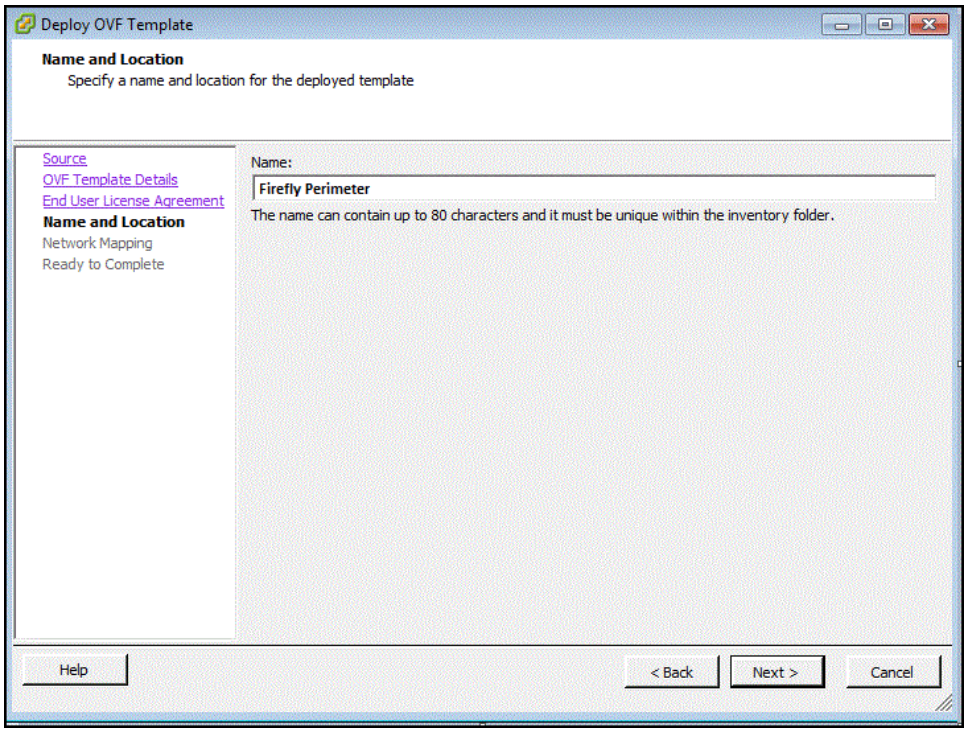

9. On the Datastore window, do not change the default settings for:

   - Datastore
   - Available Space

   Table 11 on page 35 lists the two disk formats available to store the virtual disk. You can choose one of the three options listed.

NOTE: **For detailed information on the two types of disk formats, see** http://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vsphere.vm_admin.doc_50%2FGUID-4C0F4D73-82F2-4B81-8AA7-1DD752A8A5AC.html.

Table 11: Disk Formats for Virtual Disk Storage

Disk Format	Utility
Thick Provision Lazy Zeroed	A virtual disk provisioning policy where disk space is assigned to the virtual disk when the virtual disk is created. Previously stored data is not erased when the disk space is created. The previous data is erased when the VM is used for the first time.

Table 11: Disk Formats for Virtual Disk Storage *(continued)*

Disk Format	Utility
Thick Provision Eager Zeroed	A virtual disk provisioning policy that erases the previously stored data completely and then allocates the disk space to the virtual disk. Creation of disks in this format is time consuming.

10. Select your destination network from the list and click **Next**.

11. Click **Finish** to complete the installation. See Figure 4 on page 36.

Figure 4: Firefly Perimeter Deployment

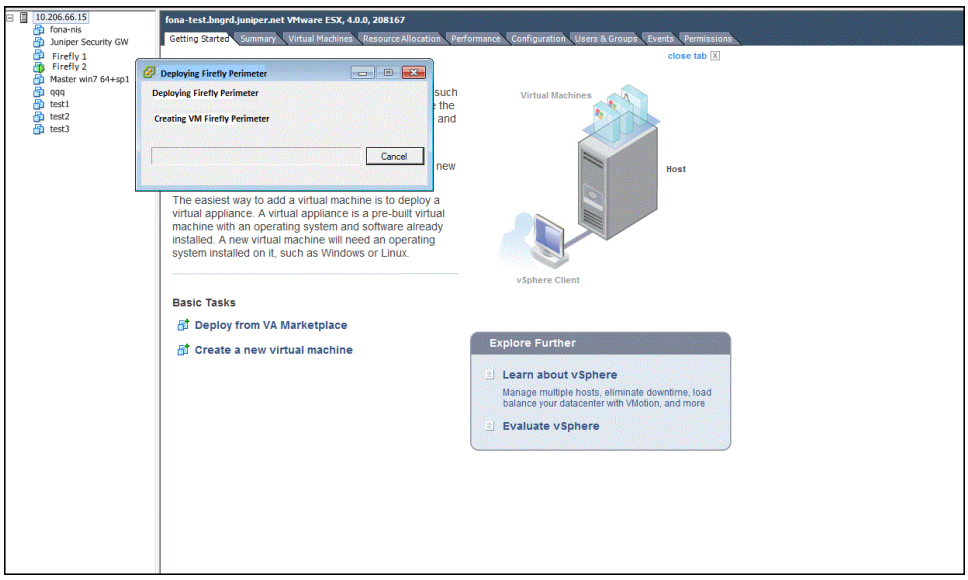

NOTE: The default Firefly Perimeter VM login ID is root with no password. By default, Firefly Perimeter is assigned a DHCP-based IP address if a DHCP server is available on the network.

Related Documentation

- Connecting the Management Device on page 36
- Powering On/Off the Device on page 37

Connecting the Management Device

- Verify That the Management Device Acquires an IP Address on page 36

Verify That the Management Device Acquires an IP Address

After you connect the management device to the services gateway, the VMware vSphere Client console automatically assigns an IP address to the management device.

NOTE: When Firefly Perimeter is powered on for the first time, it uses the factory default configuration.

Related Documentation

- Understanding Firefly Perimeter on page 3
- Installing Firefly Perimeter with VMware vSphere Client on page 33
- Powering On/Off the Device on page 37

Powering On/Off the Device

To power on Firefly Perimeter after deployment:

Click the green **Power On** icon on the menu bar.

To power off Firefly Perimeter after deployment:

Click the **Console** tab on the VMware vSphere Client interface and run the following command:

request system halt or **request system power-off**

Wait for a few minutes and then click the **Shut Down Guest** icon on the menu bar.

Related Documentation

- Understanding Firefly Perimeter on page 3
- Installing Firefly Perimeter with VMware vSphere Client on page 33
- Connecting the Management Device on page 36

PART 3

Configuration

CHAPTER 4

Firefly Perimeter Configurations

Firefly Perimeter Configuration Using the J-Web Interface

Accessing the J-Web Interface and Configuring Firefly Perimeter

To configure Firefly Perimeter using the J-Web Interface:

1. Launch a Web browser from the management device.
2. Enter the Firefly Perimeter interface IP address in the Address box.
3. Specify the default username as root. Do not enter a value in the Password box.
4. Click **Log In**. The J-Web Setup Wizard page opens. See Figure 5 on page 42.

Figure 5: J-Web Setup Wizard Page

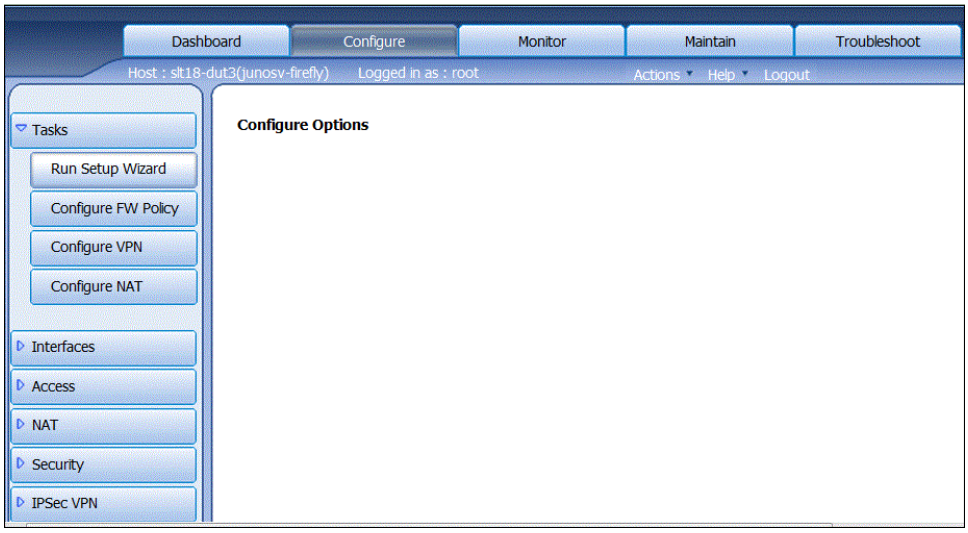

5. Click **Tasks** > **Run Setup Wizard**.

   You can use the Setup Wizard to configure a device or edit an existing configuration. See Figure 6 on page 42.

Figure 6: J-Web Configuration Page

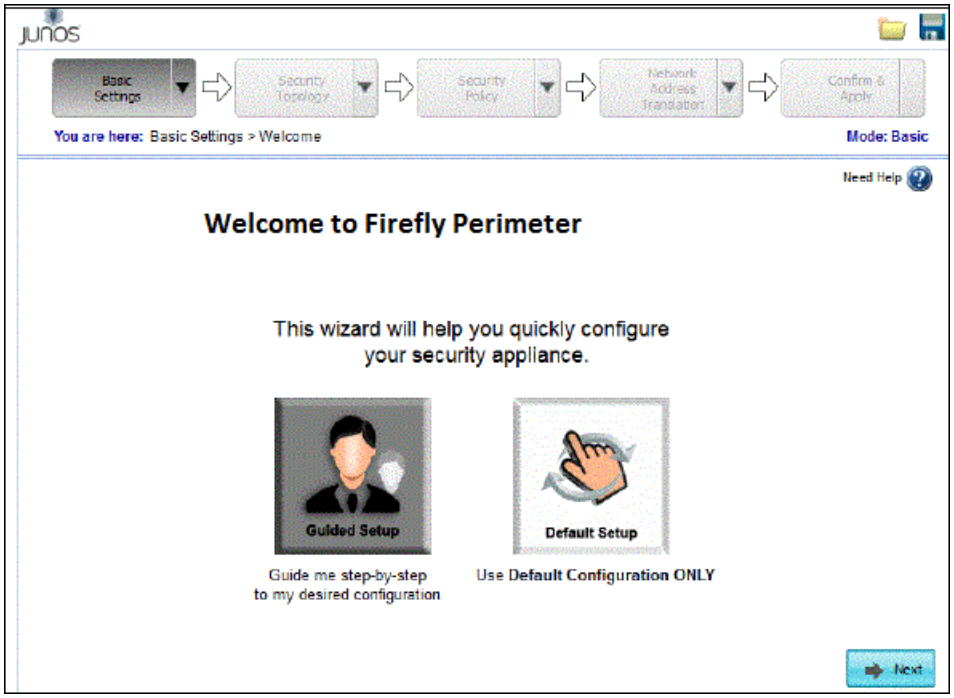

- Select the **Edit Existing Configuration** if you have already configured the wizard using the factory mode.
- Select the **Create New Configuration** to configure a device using the wizard.

Two configuration options are available:

- To enable basic options

  Select **Basic** to enable basic options. In Basic mode, you configure the device name and user account information as shown in Table 12 on page 43.

  - Device name and user account information

Table 12: Device Name and User Account Information

Field	Description
Device name	Type the name of the device. For example: **Firefly Perimeter**.
Root password	Create a default root user password.
Verify password	Verify the default root user password.
Operator	Add an administrative account in addition to the root account, which is optional. User role options include: • **Super User**: This user has full system administration rights and can add, modify, and delete settings and users. • **Operator**: This user can perform system operations such as a system reset but cannot change the configuration or add or modify users. • **Read only**: This user can only access the system and view the configuration. • **Disabled**: This user cannot access the system.

- Select either **Time Server** or **Manual**. Table 13 on page 43 lists the system time options.

Table 13: System Time Options

Field	Description
Time Server	
Host Name	Type the hostname of the time server. For example: **us.ntp.pool.org**
IP	Type the IP address of the time server in the IP address entry field. For example: 1**92.168.1.254**.
NOTE: You can either enter the hostname or the IP address.	
Manual	
Date	Click the current date in the calendar.
Time	Set the hour, minute, and seconds. Choose **AM** or **PM**.
Time Zone (mandatory)	

Table 13: System Time Options *(continued)*

Field	Description
Time Zone	Select the time zone from the list. For example: GMT Greenwich Mean Time GMT.

- To enable Advanced options:

  Select **Expert** to configure the basic options as well as the following advanced options:

  - Four or more internal zones
  - Internal zone services
  - Application of security policies between internal zones
  - A static IP address pool for Internet addressing
  - An inbound static IP addressing pool for NAT

  Click the **Need Help** icon available for detailed configuration information.

You see a success message after the basic configuration is complete. See Figure 7 on page 44.

Figure 7: Firefly Perimeter Configuration Summary

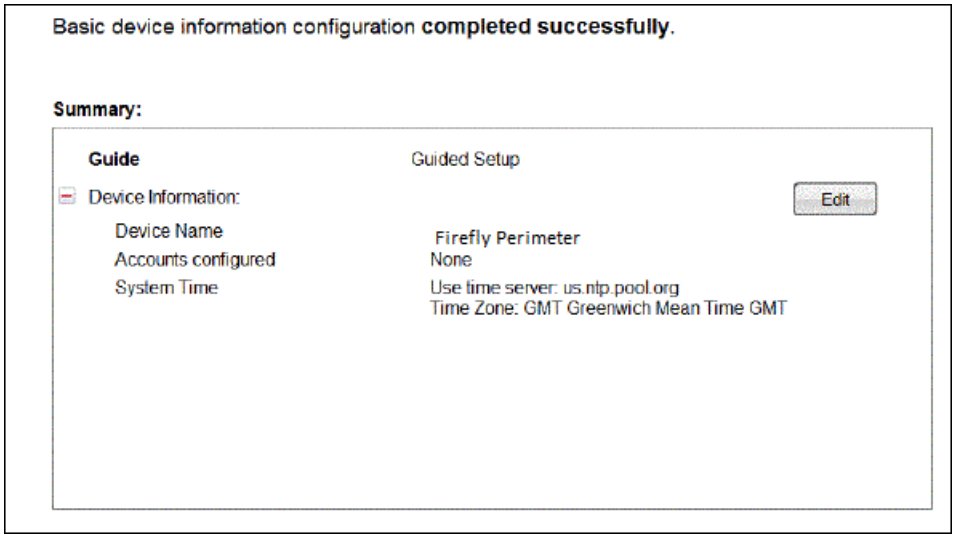

Applying the Configuration

To apply the configuration settings for Firefly Perimeter:

1. Review and ensure that the configuration settings are correct and click **Next**. The Commit Configuration page displays.
2. Click **Apply Settings** to apply the configuration changes to Firefly Perimeter.

3. Check the connectivity to Firefly Perimeter as you might lose connectivity if you have changed the management zone IP. Click the URL for reconnection instructions on how to reconnect to the device.
4. Click **Done** to complete the setup.

   After successful completion of the setup, you are redirected to the J-Web interface.

WARNING: After you complete the initial setup configuration, you can relaunch the J-Web Setup wizard by clicking Tasks > Run Setup Wizard. You can either edit an existing configuration or create a new configuration. If you decide to create a new configuration, then all the current configuration in Firefly Perimeter will be deleted.

Related Documentation

- Firefly Perimeter Basic Settings on page 28
- Powering On/Off the Device on page 37
- Firefly Perimeter Configuration Using the CLI Interface on page 45

Firefly Perimeter Configuration Using the CLI Interface

To configure Firefly Perimeter using the CLI Interface:

1. Verify that the device is powered on.
2. Log in as the root user. There is no password.
3. Start the CLI

   ```
   root#cli
   root@>
   ```

4. Enter configuration mode.

   ```
   configure
   [edit]
   root@#
   ```

5. Set the root authentication password by entering a cleartext password, an encrypted password, or an SSH public key string (DSA or RSA).

   ```
   [edit]
   root@# set system root-authentication plain-text-password
   New password: password
   Retype new password: password
   ```

6. Configure an administrative account on the device.

   ```
   [edit]
   root@# set system login user admin class super-user authentication
     plain-text-password
   ```

7. Commit the configuration to activate it on the device.

   ```
   [edit]
   ```

```
root@# commit
```

8. Login as the administrative user you configured in Step 6.
9. Configure the name of the device. If the name includes spaces, enclose the name in quotation marks (“ ”).

```
configure
[edit]
admin@# set system host-name host-name
```

10. Configure the traffic interface.

```
[edit]
admin@# set interfaces ge-0/0/1 unit 0 family inet address address/prefix-length
```

11. Configure the default route.

```
[edit]
admin@# set routing-options static route 0.0.0.0/0 next-hop gateway
```

12. Configure basic security zones and bind them to traffic interfaces.

```
[edit]
admin@# set security zones security-zone untrust interfaces ge-0/0/1
```

13. Verify the configuration.

```
[edit]
admin@# commit check
configuration check succeeds
```

14. Commit the configuration to activate it on the device.

```
[edit]
admin@# commit
commit complete
```

15. Optionally, display the configuration to verify that it is correct.

```
[edit]
user@host# show
system {
   host-name devicea;
   domain-name lab.device.net;
   domain-search [ lab.device.net device.net ];
   backup-device ip
time-zone America/Los_Angeles;
   root-authentication {
      ssh-rsa "ssh-rsa AAAAB3Nza...D9Y2gXF9ac==root@devicea.lab.device.net";
   }
name-server {
   ip
}
services {
}
   ntp {
      server ip
   }
}
interfaces {
```

```
ge-0/0/0 {
    unit 0 {
        family inet {
            address ip
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address ip
        }
    }
}
}
```

16. Commit the configuration to activate it on the device.

```
[edit]
admin@# commit
```

17. Optionally, configure more properties by adding the necessary configuration statements. Then commit the changes to activate them on the device.

```
[edit]
admin@host# commit
```

18. When you have finished configuring the device, exit configuration mode.

```
[edit]
admin@host# exit
admin@host>
```

NOTE: For additional configuration details, see:

http://www.juniper.net/techpubs/en_US/junos12.1/information-products/pathway-pages/security/security-swconfig-initial-device-config.html#configuration

Related Documentation

- Firefly Perimeter Basic Settings on page 28
- Powering On/Off the Device on page 37
- Firefly Perimeter Configuration Using the J-Web Interface on page 41

Configuring Chassis Cluster for Firefly Perimeter

- Chassis Cluster Overview on page 48
- Understanding Chassis Cluster Formation on page 49
- Chassis Cluster Quick Setup on page 49
- Configuring Chassis Cluster on page 52

- Firefly Chassis Cluster Configuration on VMware on page 58
- Deploying Firefly Perimeter Chassis Cluster Nodes at Different ESXi Hosts Using dvSwicth on page 64

Chassis Cluster Overview

Chassis clustering provides network node redundancy by grouping a pair of the same kind of Firefly Perimeter instances into a cluster. The devices must be running the same version of the Junos OS. The control ports on the respective nodes are connected to form a control plane that synchronizes the configuration and kernel state to facilitate the high availability of interfaces and services. Similarly, the data plane on the respective nodes is connected over the fabric ports to form a unified data plane. The fabric link allows for the management of cross-node flow processing and for the management of session redundancy.

The control plane software operates in active or backup mode. When configured as a chassis cluster, the two nodes back up each other, with one node acting as the primary device and the other as the secondary device, ensuring stateful failover of processes and services in the event of a system or hardware failure. If the primary device fails, the secondary device takes over processing of traffic.

The data plane software operates in active/active mode. In a chassis cluster, session information is updated as traffic traverses either device, and this information is transmitted between the nodes over the fabric link to guarantee that established sessions are not dropped when a failover occurs. In active/active mode, it is possible for traffic to ingress the cluster on one node and egress from the other node.

Chassis cluster functionality includes:

- Resilient system architecture, with a single active control plane for the entire cluster and multiple Packet Forwarding Engines. This architecture presents a single device view of the cluster.
- Synchronization of configuration and dynamic runtime states between nodes within a cluster.
- Monitoring of physical interfaces, and failover if the failure parameters cross a configured threshold.
- Support for generic routing encapsulation (GRE) and IP-over-IP (IP-IP) tunnels used to route encapsulated IPv4 or IPv6 traffic by means of two internal interfaces, gr-0/0/0 and ip-0/0/0, respectively. These interfaces are created by Junos OS at system bootup and are used only for processing GRE and IP-IP tunnels.

At any given instant, a cluster can be in one of the following states: hold, primary, secondary-hold, secondary, ineligible, and disabled. A state transition can be triggered because of any event, such as interface monitoring, Services Processing Unit (SPU) monitoring, failures, and manual failovers.

For additional information, see:

Interfaces for Security Devices

Understanding Chassis Cluster Formation

You create two Firefly Perimeter instances to form a chassis cluster, and then you set the cluster ID and node ID on each instance to join the cluster. When a device joins a cluster, it becomes a node of that cluster. With the exception of unique node settings and management IP addresses, nodes in a cluster share the same configuration.

You can deploy up to 15 chassis clusters in a Layer 2 domain. Clusters and nodes are identified in the following ways:

- A cluster is identified by a *cluster ID* specified as a number from 1 to 15.
- A cluster node is identified by a *node ID* specified as a number from 0 to 1.

Generally, on SRX Series devices, the cluster ID and node ID are written into EEPROM. However, the Firefly Perimeter VM does not emulate it. A location (boot/loader.conf) is required to save the IDs and read it out during initialization. Then the whole system (including BSD kernel) can know it is working in chassis cluster mode and does related initializations for chassis cluster.

The chassis cluster formation commands for node 0 and node 1 are as follows:

- user@host**set chassis cluster cluster-id 1 node 0 reboot**
- user@host**set chassis cluster cluster-id 1 node 1 reboot**

For additional information on chassis cluster, see:

http://www.juniper.net/techpubs/en_US/junos12.1x45/information-products/pathway-pages/security/security-chassis-cluster.html

Chassis Cluster Quick Setup

You can use the J-Web interface to set up chassis cluster for both the Firefly Perimeter devices forming a cluster.

To set up chassis cluster:

1. Launch a Web browser from the management device.
2. Enter the Firefly interface IP address in the Address box.
3. Specify the default username as root. Do not enter a value in the Password box.
4. Click **Log In**. The J-Web Setup Wizard page opens.
5. Select **Configure>Chassis Cluster>Setup**. The Chassis Cluster Setup configuration page appears.Table 14 on page 51 explains the contents of this page.
6. Configure chassis cluster using the options described in Table 14 on page 51.
7. Click **Enable** to enable chassis cluster mode on the node.
8. Select one of the following options:
   - **Enable and Reboot**: Enables chassis cluster mode and reboots the node.

A confirmation message says **Successfully enabled chassis cluster. Going to reboot now.**

Click **OK**.

- **Enable and No Reboot**: Enables chassis cluster mode without rebooting the node.

  A confirmation message is displayed.

  Click **OK**.

- **Cancel**: Cancels your entries and returns to the main configuration page.

9. Click **Reset** to reset your entries to their original values or click **Disable** to disable chassis cluster mode on the node.

Table 14: Add Chassis Cluster Setup Configuration Details

Field	Function	Action
Cluster ID	Specifies the number by which a cluster is identified.	Enter a number from 0 through 15.
Node		
Node ID	Specifies the number by which a node is identified.	Enter a number from 0 through 1.
Node Management IP Address (fxp0.0)	Specifies the management IP address of a node.	Enter a valid IP address for the management interface.
Control Link		
FPC	Specifies the FPC control link.	Select the FPC number from the list.
Port	Specifies the port to configure for the control link.	Enter a number from 0 through 2.

NOTE: For detailed information on various options used for chassis cluster see:

http://www.juniper.net/techpubs/en_US/junos12.1x45/information-products/pathway-pages/security/security-chassis-cluster.html

Configuring Chassis Cluster

You can use J-Web interface to configure the primary Firefly device.

Select **Configure>Chassis Cluster>Cluster Configuration**. The Chassis Cluster configuration page appears.

See Table 15 on page 52 for the actions available on the Chassis Cluster configuration page.

Table 16 on page 53 explains the contents of the configuration page.

See Table 17 on page 54 for Node Setting configuration details.

Table 15: Chassis Cluster Configuration Page Actions

Action	Description
Add	Adds a new or duplicate chassis cluster configuration. Enter information as specified in Table 18 on page 57.
Edit	Edits the selected chassis cluster configuration. Enter information as specified in Table 18 on page 57.
Delete	Deletes the selected chassis cluster configuration.
Actions & Commit	Commits the configuration and returns to the main configuration page.
Cancel	Cancels your entries and returns to the main configuration page.

Table 16: Chassis Cluster Configuration Page

Field	Function
Node Settings	
Node ID	Displays the node ID.
Cluster ID	Displays the cluster ID configured for the node.
Host Name	Displays the name of the node.
Backup Router	Displays the IP address used while booting.
Management Interface	Displays the management interface of the node.
IP Address	Displays the management IP address of the node.
Status	Displays the state of the redundancy group. • **Primary**–Redundancy group is active. • **Secondary**–Redundancy group is passive.
Chassis Cluster> Cluster Settings>Interfaces	
Name	Displays the physical interface name.
Member Interfaces/IP Address	Displays the member interface name or IP address configured for an interface.
Redundancy Group	Displays the redundancy group.
Chassis Cluster> Cluster Settings>Redundancy Group	
Group	Displays the redundancy group identification number.
Preempt	Displays the selected Preempt option. • **True**–Mastership can be preempted based on priority. • **False**–Mastership cannot be preempt based on priority.
Gratuitous ARP Count	Displays the number of gratuitous ARP requests that a newly elected primary device in a chassis cluster sends out to announce its presence to the other network devices.
Node Priority	Displays the assigned priority for the redundancy group on that node. The eligible node with the highest priority is elected as primary for the redundant group.

Table 17: Add Node Setting Configuration Details

Field	Function	Action
Fabric Link > Fabric Link 0 (fab0)		
Interface	Specifies fabric link 0.	Enter the interface IP fabric link 0.
Add	Adds fabric interface 0.	Click **Add**.
Delete	Deletes fabric interface 0.	Click **Delete**.
Fabric Link > Fabric Link 1 (fab1)		
Interface	Specifies fabric link 1.	Enter the interface IP for fabric link 1.
Add	Adds fabric interface 1.	Click **Add**.
Delete	Deletes fabric interface 1.	Click **Delete**.
Redundant Ethernet		
Interface	Specifies a logical interface consisting of two physical Ethernet interfaces, one on each chassis.	Enter the logical interface.
IP	Specifies redundant Ethernet IP address.	Enter redundant Ethernet IP address.
Redundancy Group	Specifies redundancy group ID number in the chassis cluster.	Select a redundancy group from the list.
Add	Adds redundant Ethernet IP address.	Click **Add**.
Delete	Deletes redundant Ethernet IP address.	Click **Delete**.
Add Redundancy Group		
Redundancy Group	Specifies the redundancy group name.	Enter the redundancy group name.
Allow preemption of primaryship	Allows a node with a better priority to initiate a failover for a redundancy group. NOTE: By default, this feature is disabled. When disabled, a node with a better priority does not initiate a redundancy group failover (unless some other factor, such as faulty network connectivity identified for monitored interfaces, causes a failover).	-

Table 17: Add Node Setting Configuration Details *(continued)*

Field	Function	Action
Gratuitous ARP Count	Specifies the number of gratuitous Address Resolution Protocol requests that a newly elected master sends out on the active redundant Ethernet interface child links to notify network devices of a change in mastership on the redundant Ethernet interface links.	Enter a value from 1 to 16. The default is 4.
node0 priority	Specifies the priority value of node0 for a redundancy group.	Enter the node priority number as 0.
node1 priority	Specifies the priority value of node1 for a redundancy group.	Select the node priority number as 1.
Interface Monitor		
Interface	Specifies the number of redundant Ethernet interfaces to be created for the cluster.	Select the interface from the list.
Weight	Specifies the weight for the interface to be monitored.	Enter a value from 1 to 125..
Add	Adds interfaces to be monitored by the redundancy group and their respective weights.	Click **Add**.
Delete	Deletes interfaces to be monitored by the redundancy group along with their respective weights.	Select the interface from the configured list and click **Delete**.
IP Monitoring		
Weight	Specifies the global threshold for IP monitoring.	Enter a value from 0 to 255.
Threshold	Specifies the global threshold for IP monitoring.	Enter a value from 0 to 255.
Retry Count	Specifies the number of retries needed to declare reachability failure.	Enter a value from 5 to 15.
Retry Interval	Specifies the time interval in seconds between retries.	Enter a value from 1 to 30.
IPV4 Addresses to be monitored		
IP	Specifies the IPv4 addresses to be monitored for reachability.	Enter the IPv4 addresses.

Table 17: Add Node Setting Configuration Details *(continued)*

Field	Function	Action
Weight	Specifies the weight for the redundancy group interface to be monitored.	Enter the weight.
Interface	Specifies the logical interface through which to monitor this IP address.	Enter the logical interface address.
Secondary IP address	Specifies the source address for monitoring packets on a secondary link.	Enter the secondary IP address.
Add	Adds the IPv4 addresses to be monitored.	Click **Add**.
Delete	Delete the IPv4 addresses to be monitored.	Select the item from the list and click **Delete**.

Table 18: Edit Node Setting Configuration Details

Field	Function	Action
Node Settings		
Host Name	Specifies the name of the host.	Enter the name of the host.
Backup Router	Specifies the backup router to be used during failover.	Specifies the backup router to be used during failover.
Destination		
IP	Adds the destination address.	Click **Add**.
Delete	Deletes the destination address.	Click **Delete**.
Interface		
Interface	Specifies the interfaces available for the router. NOTE: Allows you to add and edit two interfaces for each fabric link.	Select an option.
IP	Specifies the interface IP address.	Enter the interface IP address.
Add	Adds the interface.	Click **Add**.
Delete	Deletes the interface.	Click **Delete**.

Firefly Chassis Cluster Configuration on VMware

This topic provides information on Firefly Chassis Cluster Configuration using the VMware vSphere Client. This topic explains how to connect control interface via control vSwitch, fabric interface via fabric vSwitch and data interface via data vSwitch. Make sure that chassis cluster is set up for both the Firefly VMs.

For configuration examples on chassis cluster using CLI, see:

http://www.juniper.net/techpubs/en_US/junos12.1x45/information-products/pathway-pages/security/security-chassis-cluster.html.

- Connecting Control Interface via Control vSwitch Using the VMware vSphere Client on page 58
- Connecting Fabric Interface via Fabric vSwitch Using the VMware vSphere Client on page 61
- Connecting Data Interface via Data vSwitch Using the VMware vSphere Client on page 63

Connecting Control Interface via Control vSwitch Using the VMware vSphere Client

To connect the control interface via control switch:

1. Choose **Configuration**->**Networking**.
2. Click **Add Networking** to create a vSwitch for control link.

   Choose the following attributes:

   - Connection Type
     - Virtual Machines
   - Network Access
     - Create a vSphere stand switch
     - No physical adapters
   - Port Group Properties
     - Network Label: chassis cluster Control
     - VLAN ID: None(0)

NOTE:

Port group are not VLAN. It does not segment the vSwitch into separate broadcast domains unless they have different VLAN tags.

- If dedicated vSwitch, you can use the default VLAN tag (0) or specify a VLAN tag.
- If shared vSwitch and using port group, you must assign a VLAN tag on that port group for each HA links.

After creating the control vSwitch, you can use the vSwitch default settings.

3. Click **Edit Settings** of both Firefly VMs to add the control interface (Network adapter 2) into control vSwitch.

See Figure 8 on page 59 for vSwitch 1 Properties and Figure 9 on page 60 for Virtual Machine Properties for Control vSwitch.

Figure 8: vSwitch 1 Properties

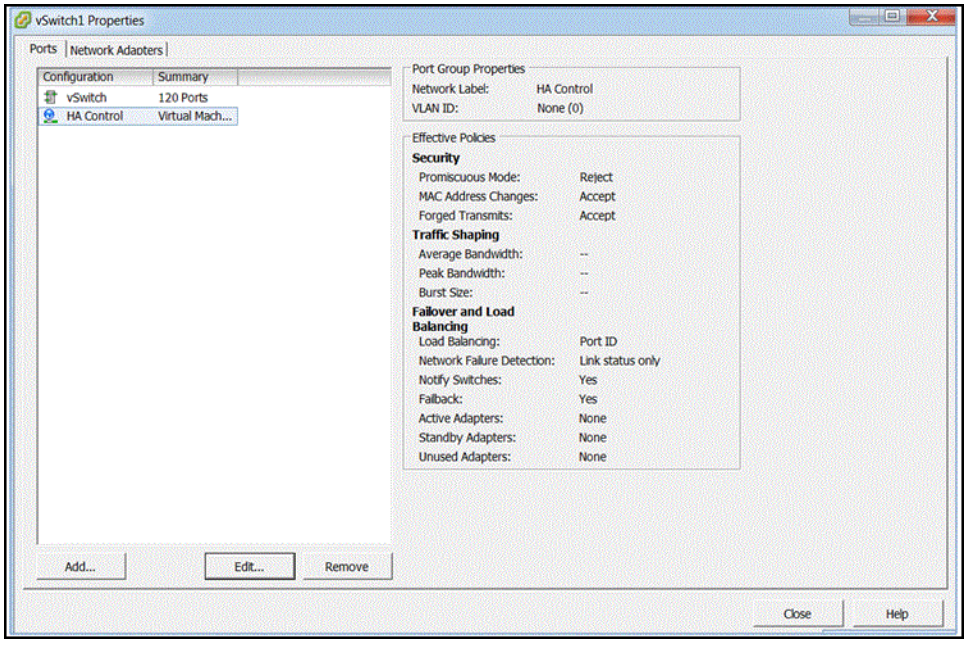

Figure 9: Virtual Machine Properties for Control vSwitch

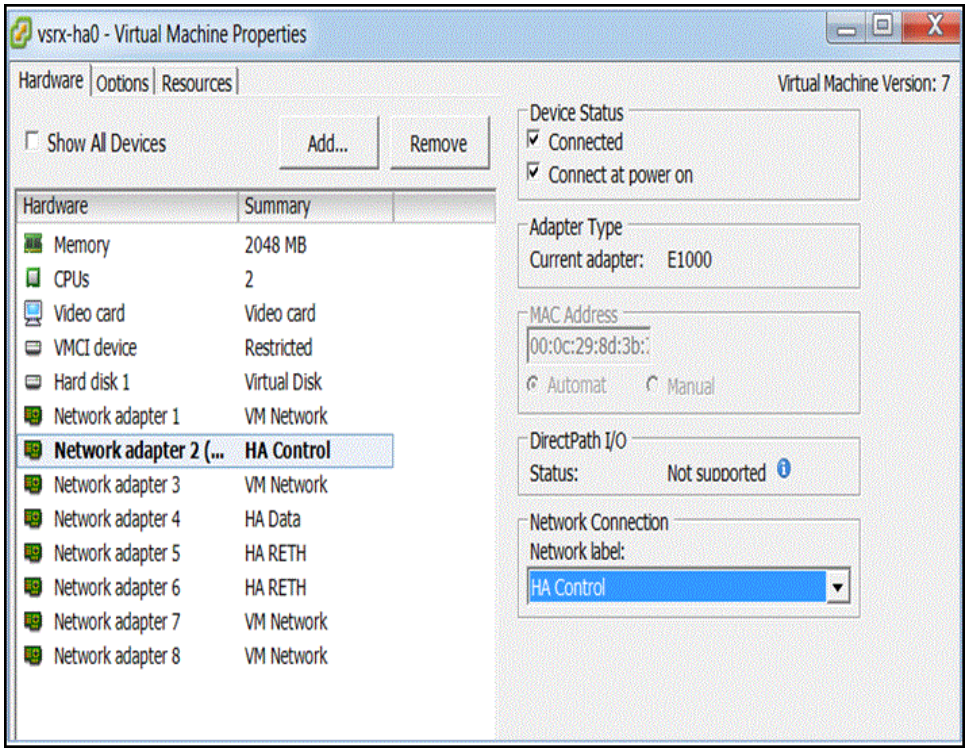

The control interface will hence be connected via the control vSwitch using the above procedure.

Connecting Fabric Interface via Fabric vSwitch Using the VMware vSphere Client

1. Choose **Configuration**->**Networking**.
2. Click **Add Networking** to create a vSwitch for fabric link.

   Choose the following attributes:

   - Connection Type
     - Virtual Machines
   - Network Access
     - Create a vSphere stand switch
     - No physical adapters
   - Port Group Properties
     - Network Label: chassis cluster Fabric
     - VLAN ID: None(0)

   NOTE:
   Port group are not VLAN. It does not segment the vSwitch into separate broadcast domains unless they have different VLAN tags.

   - If dedicated vSwitch, you can use the default VLAN tag (0) or specify a VLAN tag.
   - If shared vSwitch and using port group, you must assign a VLAN tag on that port group for each HA links.

   Click on **Properties** to turn on the following features:

   - **General**-> **Advanced Properties**:
     - MTU: 9000
   - **Security**-> **Effective Polices**:
     - MAC Address Changes: Accept
     - Forged Transmits: Accept
3. Click **Edit Settings** of both Firefly VMs to add the fabric interface into fabric vSwitch.

   NOTE: Network adaptor 4 is used in this example, which is configurable in Junos.

See Figure 10 on page 62 for vSwitch 2 Properties and Figure 11 on page 62 for Virtual Machine Properties for Fabric vSwitch.

Figure 10: vSwitch 2 Properties

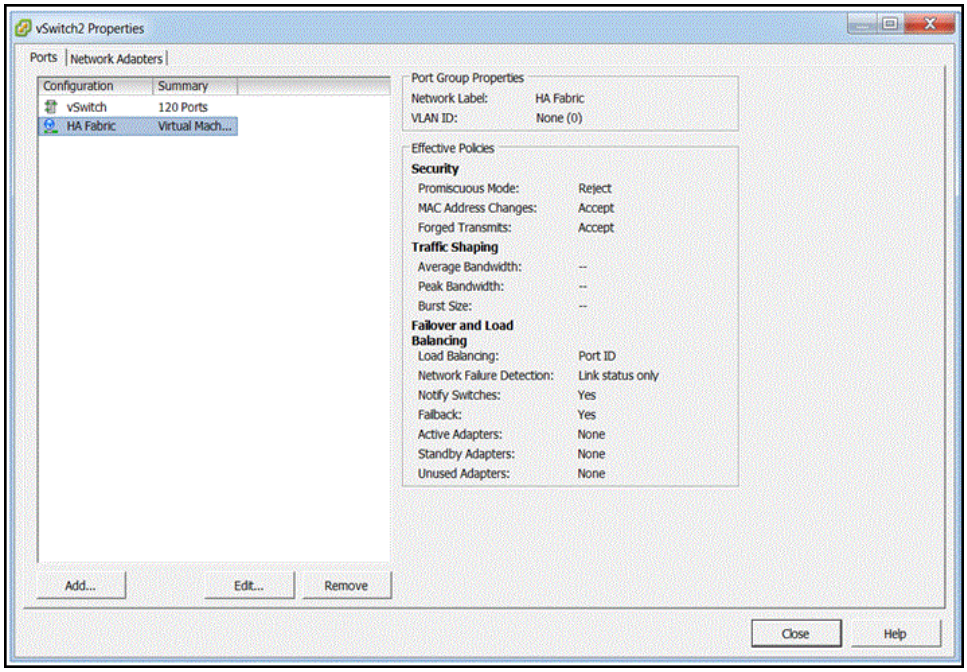

Figure 11: Virtual Machine Properties for Fabric vSwitch

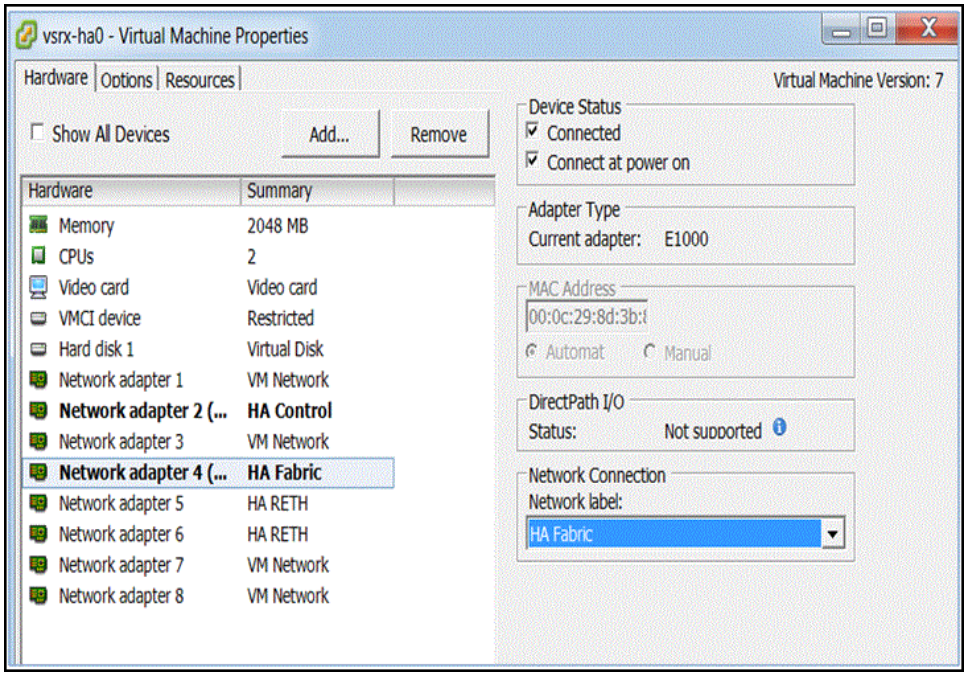

The fabric interface will hence be connected via the fabric vSwitch using the above procedure.

Connecting Data Interface via Data vSwitch Using the VMware vSphere Client

Add all the redundant interfaces into data traffic vSwitch like standalone mode.

1. Choose **Configuration**->**Networking**.
2. Click **Add Networking** to create a vSwitch for fabric link.

   Choose the following attributes:

   - Connection Type
     - Virtual Machines
   - Network Access
     - Create a vSphere stand switch
     - No physical adapters
   - Port Group Properties
     - Network Label: chassis cluster Reth
     - VLAN ID: None(0)

     Click on **Properties** to turn on the following features:

     - **General**-> **Advanced Properties**:
       - MTU: 9000
     - **Security**-> **Effective Polices**:
       - MAC Address Changes: Accept
       - Forged Transmits: Accept

See Figure 12 on page 64 for Virtual Machine Properties for Data vSwicth.

Figure 12: Virtual Machine Properties for Data vSwitch

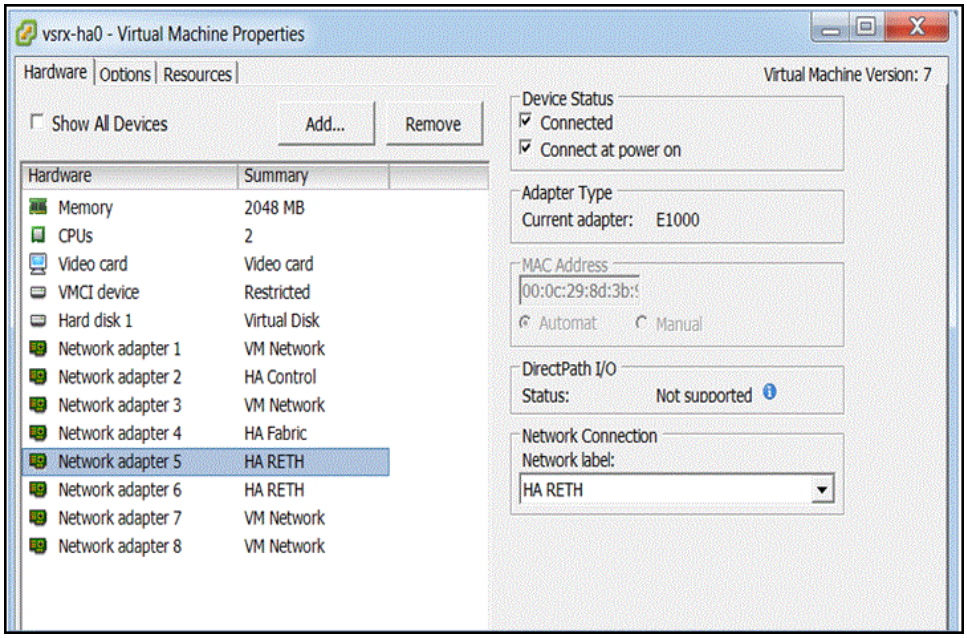

The data interface will hence be connected via the data vSwitch using the above procedure.

Deploying Firefly Perimeter Chassis Cluster Nodes at Different ESXi Hosts Using dvSwicth

In this method, we use the private vlan (PVLAN) feature of dvSwitch. There is no need to change the external switch configurations.

On the VMware vSphere Client, for dvSwitch, there are two private VLAN IDs, the primary private VLAN ID and the secondary private VLAN ID.

Select **Community** in the drop down menu for secondary VLAN ID type.

Use the two secondary private VLAN IDs for Firefly Perimeter control and fabric link. See Figure 13 on page 65 and Figure 14 on page 65.

Figure 13: dvPortGroup3 Settings

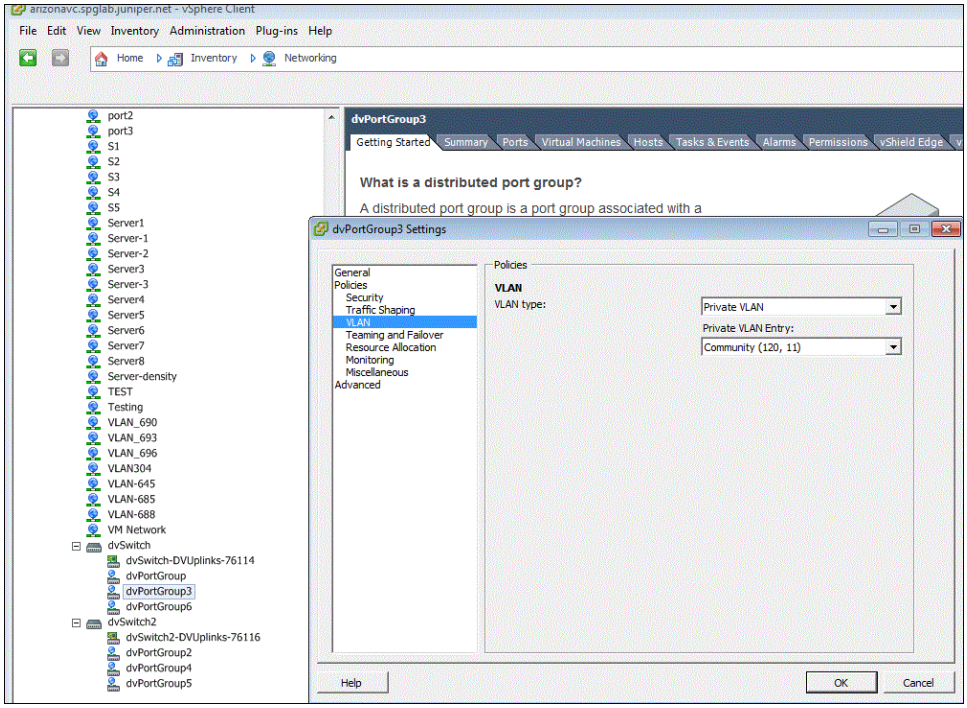

Figure 14: dvPortGroup6 Settings

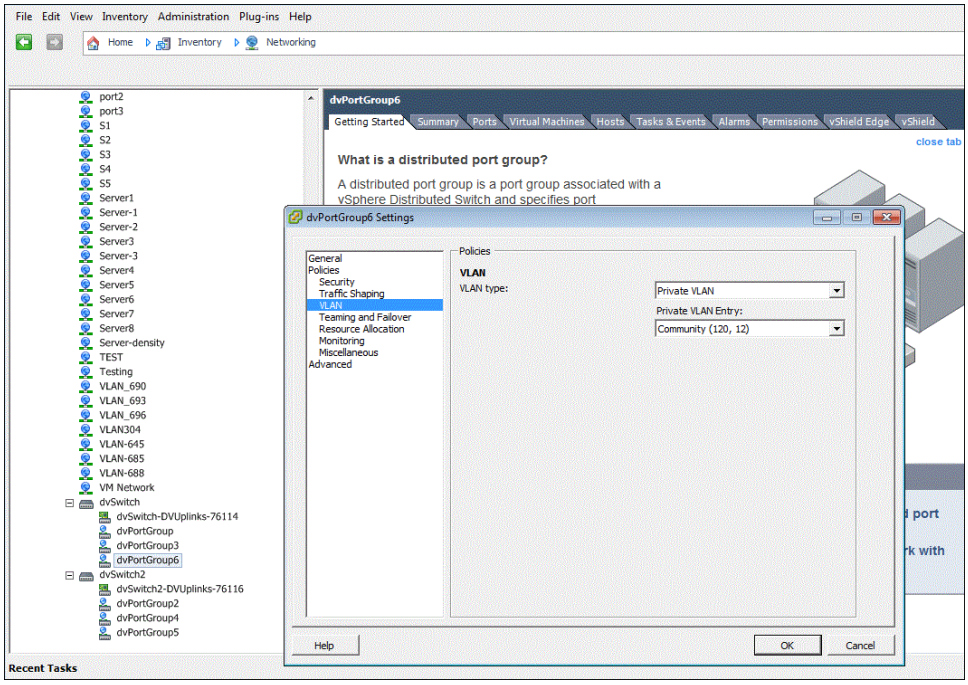

NOTE: Configurations above are required at external switch, to which distributed switch uplinks are connected. If the link at external switch has native vlan, then distributed switch port group config can have vlan as none. Otherwise, vlan should be used.

You can also use regular VLAN on a distributed switch to deploy Firefly Perimeter chassis cluster nodes at different ESXi Hosts using dvSwitch. Regular VLAN works similar to a physical switch. If you want to use regular VLAN instead of PVLAN, disable IGMP snooping for chassis cluster links.

However, use of PVLAN is recommended because:

- PVLAN does not impose IGMP snooping.
- PVLAN can save VLAN IDs.

Related Documentation

- Understanding Firefly Perimeter on page 3
- Features Supported on Firefly Perimeter with VMware on page 3

PART 4

Index

Index