



QUANTUM

16 July 2025

SITE TO SITE VPN

R81.10

Administration Guide



Check Point Copyright Notice

© 2021 - 2025 Check Point Software Technologies Ltd.

All rights reserved. This product and related documentation are protected by copyright and distributed under licensing restricting their use, copying, distribution, and decompilation. No part of this product or related documentation may be reproduced in any form or by any means without prior written authorization of Check Point. While every precaution has been taken in the preparation of this book, Check Point assumes no responsibility for errors or omissions. This publication and features described herein are subject to change without notice.

RESTRICTED RIGHTS LEGEND:

Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

TRADEMARKS:

Refer to the [Copyright page](#) for a list of our trademarks.

Refer to the [Third Party copyright notices](#) for a list of relevant copyrights and third-party licenses.

Important Information



Latest Software

We recommend that you install the most recent software release to stay up-to-date with the latest functional improvements, stability fixes, security enhancements and protection against new and evolving attacks.



Certifications

For third party independent certification of Check Point products, see the [Check Point Certifications page](#).



Check Point R81.10

For more about this release, see the R81.10 [home page](#).



Latest Version of this Document in English

Open the latest version of this [document in a Web browser](#).
Download the latest version of this [document in PDF format](#).



Feedback

Check Point is engaged in a continuous effort to improve its documentation.
[Please help us by sending your comments](#).

Revision History

Date	Description
19 May 2025	Updated: ▪ "Dead Peer Detection" on page 147
23 October 2023	Updated: ▪ "Tunnel Management" on page 146
25 June 2023	Updated: ▪ "Basic Site to Site VPN Configuration" on page 45
05 April 2023	Updated: ▪ "Route Injection Mechanism" on page 154
06 February 2023	Updated: ▪ "vpn debug" on page 217
16 January 2023	Updated: ▪ "Getting Started with Site-to-Site VPN" on page 32 ▪ "Multiple Entry Point (MEP) VPNs" on page 174 ▪ "Resolving Connectivity Issues" on page 194
11 September 2022	Updated: ▪ "Multiple Entry Point (MEP) VPNs" on page 174
07 July 2022	Updated : ▪ "Perfect Forward Secrecy" on page 59
14 June 2022	In the HTML version, added glossary terms in the text.
27 March 2022	Updated: ▪ "Getting Started with Site-to-Site VPN" on page 32 - added a note that the Encryption Domain Per Community feature requires Security Gateways R80.40 and higher ▪ "Perfect Forward Secrecy" on page 59 - added notes
24 November 2021	Updated: ▪ "Tunnel Management" on page 146 - changed the procedure to enable or disable the feature "Delete IKE SAs for dead peer"

Date	Description
07 November 2021	Updated: ▪ "Basic Site to Site VPN Configuration" on page 45 - added a note that in Star community, the "Center Gateways" section does not support Quantum Spark appliances
02 August 2021	Updated: ▪ "Basic Site to Site VPN Configuration" on page 45 ▪ "Route Injection Mechanism" on page 154
06 July 2021	First release of this document

Table of Contents

Glossary	15
Check Point VPN	25
IPsec VPN	25
VPN Components	25
Understanding the Terminology	26
Site-to-Site VPN	27
Sample Site-to-Site VPN Deployment	27
VPN Communities	28
Sample Combination VPN Community	29
Routing VPN Traffic	29
Granular Routing Control	30
IPv6 Support and Limitations	30
Getting Started with Site-to-Site VPN	32
Basic Site to Site VPN Configuration	45
Configuring a Star or Meshed Community Between Internally Managed Security Gateways	45
Configuring a VPN with External Security Gateways Using Certificates	48
Configuring a VPN with External Security Gateways Using Pre-Shared Secret	50
Firewall Control Connections in VPN Communities	52
Why Turning off Implied Rules Blocks Firewall Control Connections	52
Allowing Firewall Control Connections Inside a VPN	53
Discovering Which Services are Used for Control Connections	53
Simplified and Traditional Modes	54
IPsec and IKE	55
Overview	55
IKE Phase I	55
IKE Phase II (Quick mode or IPSec Phase)	56
IKEv1 and IKEv2	56

Methods of Encryption and Integrity	56
Diffie Hellman Groups	57
Phase I modes	58
Renegotiating IKE & IPsec Lifetimes	58
Perfect Forward Secrecy	59
IP Compression	60
Subnets and Security Associations	60
Unique SA Per Pair of Peers	60
IKE DoS Protection	61
Understanding DoS Attacks	61
IKE DoS Attacks	61
Defense Against IKE DoS Attacks	62
SmartConsole IKE DoS Attack Protection Settings	62
Advanced IKE DoS Attack Protection Settings	63
Protection After Successful Authentication	64
Client Properties	65
Configuring Advanced IKE Properties	65
VPN Community Object - Encryption Settings	65
VPN Community Object - Advanced Settings	67
Link Selection	69
Link Selection Overview	69
Configuring IP Selection by Remote Peer	69
Last Known Available Peer IP Address	70
Configuring Outgoing Route Selection	70
When Responding to a Remotely Initiated Tunnel	71
Using Route Based Probing	71
Source IP Address Settings	72
Outgoing Link Tracking	73
Link Selection Scenarios	73
Security Gateway with a Single External Interface	73

Security Gateway with Several IP Addresses Used by Different Parties	74
Security Gateway with an Interface Behind a Static NAT Device	75
Utilizing Load Sharing	76
Load Sharing with Multiple External Interfaces on Each End	76
Load Sharing with Multiple External Interfaces on One End	77
Service Based Link Selection	79
Configuring Service Based Link Selection	79
Service Based Link Selection Scenarios	80
Service Based Link Selection with Two Interfaces on Each End	80
Service Based Link Selection with Multiple Interfaces on Each End	82
Service Based Link Selection with Two Interfaces on One End	83
Trusted Links	84
Configuring Trusted Links	85
Trusted Links Scenarios	85
Using Trusted Links with Service Based Link Selection	86
On Demand Links (ODL)	87
Configuring On Demand Links	88
Link Selection and ISP Redundancy	89
Configuring Link Selection and ISP Redundancy	89
Link Selection and ISP Redundancy	90
Link Selection with non-Check Point Devices	92
Public Key Infrastructure	94
Need for Integration with Different PKI Solutions	94
Supporting a Wide Variety of PKI Solutions	94
PKI and Remote Access Users	95
PKI Deployments and VPN	95
Simple Deployment Internal CA	95
CA of An External Security Management Server	95
CA Services Over the Internet	96
CA Located on the LAN	97

Trusting An External CA	97
Subordinate Certificate Authorities	98
Enrolling a Managed Entity	98
Validation of a Certificate	98
Revocation Checking	99
Enrolling with a Certificate Authority	99
CRL	104
OCSP	104
CRL Prefetch-Cache	104
Special Considerations for the CRL Pre-fetch Mechanism	105
CRL Grace Period	105
Special Considerations for PKI	106
Using the Internal CA vs. Deploying a Third Party CA	106
Distributed Key Management and Storage	106
Configuration of PKI Operations	106
Trusting a CA - Step-By-Step	106
Trusting an ICA	107
Trusting an Externally Managed CA	107
Trusting an OPSEC Certified CA	107
Certificate Revocation (All CA Types)	109
Certificate Recovery and Renewal	109
Recovery and Renewal with Internal CA	109
CA Certificate Rollover	109
Managing a CA Certificate Rollover	110
CA Certificate Rollover CLI	111
Adding Matching Criteria to the Validation Process	111
CRL Cache Usage	111
Modifying the CRL Pre-Fetch Cache	112
Configuring CRL Grace Period	112
Configuring OCSP	112

Domain Based VPN	113
Overview of Domain-based VPN	113
VPN Routing and Access Control	114
Configuring VPN Routing in Domain Based VPN	114
Configuring VPN Routing for Security Gateways in SmartConsole	114
Configuration in the VPN Configuration File	115
Configuring the 'Accept VPN Traffic Rule'	115
Configuring Multiple Hubs	116
Configuring VPN Routing and Access Control on Security Management Server A	116
Configuring VPN Routing and Access Control on Security Management Server B	117
VPN with One or More LSM Profiles	118
Route Based VPN	119
Overview of Route-based VPN	119
VPN Tunnel Interface (VTI)	119
Using Dynamic Routing Protocols	121
VTIs in a Clustered Environment	122
Configuring VTIs in Gaia Operating System	122
Enabling Route Based VPN	122
Configuring Numbered VTIs - Example	123
Enabling Dynamic Routing Protocols on VTIs - Example	132
Configuring Anti-Spoofing on VTIs in SmartConsole	134
Routing Multicast Packets Through VPN Tunnels	135
Large Scale VPN	136
Configuring LSV	136
Monitoring LSV Peers and Tunnels	145
Tunnel Management	146
Overview of Tunnel Management	146
Permanent Tunnels	146
Permanent Tunnels in a MEP Environment	147
Tunnel Testing for Permanent Tunnels	147

Terminating Permanent Tunnels	147
Dead Peer Detection	147
Permanent Tunnel Based on Dead Peer Detection	147
VPN Tunnel Sharing	149
Configuring Tunnel Features	149
Permanent Tunnels	150
Advanced Permanent Tunnel Configuration	151
Tracking Options	152
Route Injection Mechanism	154
Overview of Route Injection	154
Automatic RIM	154
Custom Scripts	157
Injecting Peer Security Gateway Interfaces	158
Configuring RIM	159
Configuring RIM in a Star Community	159
Configuring RIM in a Meshed Community	160
Enabling the RIM_inject_peer_interfaces flag	161
Configuring RIM in Gaia	161
Wire Mode	164
Overview of Wire Mode	164
Wire Mode Scenarios	164
Wire Mode in a MEP Configuration	164
Wire Mode with Route Based VPN	165
Wire Mode Between Two VPN Communities	166
Special Considerations for Wire Mode	167
Configuring Wire Mode	168
Enabling Wire Mode on a VPN Community	168
Enabling Wire Mode on a Specific Security Gateway	168
Directional VPN Enforcement	169
Overview of Directional VPN	169

Directional Enforcement within a Community	169
Configurable Objects in a Direction	170
Directional Enforcement between Communities	171
Configuring Directional VPN Within a Community	172
Configuring Directional VPN Between Communities	173
Multiple Entry Point (MEP) VPNs	174
Overview of MEP	174
VPN High Availability Using MEP or Clustering	174
Implementation	174
Explicit MEP	175
MEP Selection Methods	177
Tracking	182
Implicit MEP	182
Routing Return Packets	186
IP Pool NAT	186
Route Injection Mechanism	186
Special Considerations	187
Configuring MEP	187
Configuring Explicit MEP	188
Configuring Implicit MEP	188
Configuring IP Pool NAT	192
Resolving Connectivity Issues	194
IPsec NAT-Traversal	194
Configuring NAT-Traversal	194
Advanced NAT-T Configuration	194
Command Line Reference	196
Syntax Legend	197
ike debug	199
vpn	208
vpn check_ttm	211

vpn compreset	212
vpn compstat	213
vpn crl_zap	214
vpn crlview	215
vpn debug	217
vpn dll	227
vpn drv	228
vpn dump_psk	229
vpn ipafile_check	230
vpn ipafile_users_capacity	231
vpn macutil	232
vpn mep_refresh	233
vpn neo_proto	234
vpn nssm_topology	235
vpn overlap_encdom	236
vpn rim_cleanup	238
vpn rll	239
vpn set_slim_server	240
vpn set_snx_encdom_groups	241
vpn set_trac	242
vpn shell	243
vpn show_tcpt	250
vpn sw_topology	251
vpn tu	252
vpn tu del	254
vpn tu list	257
vpn tu mstats	259
vpn tu tlist	260
vpn ver	262
mcc	263

mcc add	265
mcc add2main	266
mcc del	267
mcc lca	268
mcc main2add	269
mcc show	270
Working with Kernel Parameters	272
Kernel Debug	273
Appendix	274
Configuring specific settings for each VPN Community	274

Glossary

A

Anti-Bot

Check Point Software Blade on a Security Gateway that blocks botnet behavior and communication to Command and Control (C&C) centers. Acronyms: AB, ABOT.

Anti-Spam

Check Point Software Blade on a Security Gateway that provides comprehensive protection for email inspection. Synonym: Anti-Spam & Email Security. Acronyms: AS, ASPAM.

Anti-Virus

Check Point Software Blade on a Security Gateway that uses real-time virus signatures and anomaly-based protections from ThreatCloud to detect and block malware at the Security Gateway before users are affected. Acronym: AV.

Application Control

Check Point Software Blade on a Security Gateway that allows granular control over specific web-enabled applications by using deep packet inspection. Acronym: APPI.

Audit Log

Log that contains administrator actions on a Management Server (login and logout, creation or modification of an object, installation of a policy, and so on).

B

Bridge Mode

Security Gateway or Virtual System that works as a Layer 2 bridge device for easy deployment in an existing topology.

C

Cluster

Two or more Security Gateways that work together in a redundant configuration - High Availability, or Load Sharing.

Cluster Member

Security Gateway that is part of a cluster.

Compliance

Check Point Software Blade on a Management Server to view and apply the Security Best Practices to the managed Security Gateways. This Software Blade includes a library of Check Point-defined Security Best Practices to use as a baseline for good Security Gateway and Policy configuration.

Content Awareness

Check Point Software Blade on a Security Gateway that provides data visibility and enforcement. Acronym: CTNT.

CoreXL

Performance-enhancing technology for Security Gateways on multi-core processing platforms. Multiple Check Point Firewall instances are running in parallel on multiple CPU cores.

CoreXL Firewall Instance

On a Security Gateway with CoreXL enabled, the Firewall kernel is copied multiple times. Each replicated copy, or firewall instance, runs on one processing CPU core. These firewall instances handle traffic at the same time, and each firewall instance is a complete and independent firewall inspection kernel. Synonym: CoreXL FW Instance.

CoreXL SND

Secure Network Distributer. Part of CoreXL that is responsible for: Processing incoming traffic from the network interfaces; Securely accelerating authorized packets (if SecureXL is enabled); Distributing non-accelerated packets between Firewall kernel instances (SND maintains global dispatching table, which maps connections that were assigned to CoreXL Firewall instances). Traffic distribution between CoreXL Firewall instances is statically based on Source IP addresses, Destination IP addresses, and the IP 'Protocol' type. The CoreXL SND does not really "touch" packets. The decision to stick to a particular FWK daemon is done at the first packet of connection on a very high level, before anything else. Depending on the SecureXL settings, and in most of the cases, the SecureXL can be offloading decryption calculations. However, in some other cases, such as with Route-Based VPN, it is done by FWK daemon.

CPUSE

Check Point Upgrade Service Engine for Gaia Operating System. With CPUSE, you can automatically update Check Point products for the Gaia OS, and the Gaia OS itself.

D

DAIP Gateway

Dynamically Assigned IP (DAIP) Security Gateway is a Security Gateway, on which the IP address of the external interface is assigned dynamically by the ISP.

Data Loss Prevention

Check Point Software Blade on a Security Gateway that detects and prevents the unauthorized transmission of confidential information outside the organization. Acronym: DLP.

Data Type

Classification of data in a Check Point Security Policy for the Content Awareness Software Blade.

Distributed Deployment

Configuration in which the Check Point Security Gateway and the Security Management Server products are installed on different computers.

Dynamic Object

Special object type, whose IP address is not known in advance. The Security Gateway resolves the IP address of this object in real time.

E

Encryption Domain

The networks that a Security Gateway protects and for which it encrypts and decrypts VPN traffic.

Endpoint Policy Management

Check Point Software Blade on a Management Server to manage an on-premises Harmony Endpoint Security environment.

Expert Mode

The name of the elevated command line shell that gives full system root permissions in the Check Point Gaia operating system.

G

Gaia

Check Point security operating system that combines the strengths of both SecurePlatform and IPSO operating systems.

Gaia Clish

The name of the default command line shell in Check Point Gaia operating system. This is a restricted shell (role-based administration controls the number of commands available in the shell).

Gaia Portal

Web interface for the Check Point Gaia operating system.

H

Hotfix

Software package installed on top of the current software version to fix a wrong or undesired behavior, and to add a new behavior.

HTTPS Inspection

Feature on a Security Gateway that inspects traffic encrypted by the Secure Sockets Layer (SSL) protocol for malware or suspicious patterns. Synonym: SSL Inspection. Acronyms: HTTPSI, HTTPSi.

I

ICA

Internal Certificate Authority. A component on Check Point Management Server that issues certificates for authentication.

Identity Awareness

Check Point Software Blade on a Security Gateway that enforces network access and audits data based on network location, the identity of the user, and the identity of the computer. Acronym: IDA.

Identity Logging

Check Point Software Blade on a Management Server to view Identity Logs from the managed Security Gateways with enabled Identity Awareness Software Blade.

Internal Network

Computers and resources protected by the Firewall and accessed by authenticated users.

IPS

Check Point Software Blade on a Security Gateway that inspects and analyzes packets and data for numerous types of risks (Intrusion Prevention System).

IPsec VPN

Check Point Software Blade on a Security Gateway that provides a Site to Site VPN and Remote Access VPN access.

J

Jumbo Hotfix Accumulator

Collection of hotfixes combined into a single package. Acronyms: JHA, JHF, JHFA.

K

Kerberos

An authentication server for Microsoft Windows Active Directory Federation Services (ADFS).

L

Log Server

Dedicated Check Point server that runs Check Point software to store and process logs.

Logging & Status

Check Point Software Blade on a Management Server to view Security Logs from the managed Security Gateways.

M

Management Interface

(1) Interface on a Gaia Security Gateway or Cluster member, through which Management Server connects to the Security Gateway or Cluster member. (2) Interface on Gaia computer, through which users connect to Gaia Portal or CLI.

Management Server

Check Point Single-Domain Security Management Server or a Multi-Domain Security Management Server.

Manual NAT Rules

Manual configuration of NAT rules by the administrator of the Check Point Management Server.

Mobile Access

Check Point Software Blade on a Security Gateway that provides a Remote Access VPN access for managed and unmanaged clients. Acronym: MAB.

Multi-Domain Log Server

Dedicated Check Point server that runs Check Point software to store and process logs in a Multi-Domain Security Management environment. The Multi-Domain Log Server consists of Domain Log Servers that store and process logs from Security Gateways that are managed by the corresponding Domain Management Servers. Acronym: MDLS.

Multi-Domain Server

Dedicated Check Point server that runs Check Point software to host virtual Security Management Servers called Domain Management Servers. Synonym: Multi-Domain Security Management Server. Acronym: MDS.

N

Network Object

Logical object that represents different parts of corporate topology - computers, IP addresses, traffic protocols, and so on. Administrators use these objects in Security Policies.

Network Policy Management

Check Point Software Blade on a Management Server to manage an on-premises environment with an Access Control and Threat Prevention policies.

O**Open Server**

Physical computer manufactured and distributed by a company, other than Check Point.

P**Provisioning**

Check Point Software Blade on a Management Server that manages large-scale deployments of Check Point Security Gateways using configuration profiles. Synonyms: SmartProvisioning, SmartLSM, Large-Scale Management, LSM.

Q**QoS**

Check Point Software Blade on a Security Gateway that provides policy-based traffic bandwidth management to prioritize business-critical traffic and guarantee bandwidth and control latency.

R**Remote Access VPN**

An encrypted tunnel between remote access clients (such as Endpoint Security VPN) and a Security Gateway.

Route-Based VPN

A routing method for participants in a VPN community, defined by network routes.

Rule

Set of traffic parameters and other conditions in a Rule Base (Security Policy) that cause specified actions to be taken for a communication session.

Rule Base

All rules configured in a given Security Policy. Synonym: Rulebase.

S

SecureXL

Check Point product on a Security Gateway that accelerates IPv4 and IPv6 traffic that passes through a Security Gateway.

Security Gateway

Dedicated Check Point server that runs Check Point software to inspect traffic and enforce Security Policies for connected network resources.

Security Management Server

Dedicated Check Point server that runs Check Point software to manage the objects and policies in a Check Point environment within a single management Domain. Synonym: Single-Domain Security Management Server.

Security Policy

Collection of rules that control network traffic and enforce organization guidelines for data protection and access to resources with packet inspection.

SIC

Secure Internal Communication. The Check Point proprietary mechanism with which Check Point computers that run Check Point software authenticate each other over SSL, for secure communication. This authentication is based on the certificates issued by the ICA on a Check Point Management Server.

Site to Site VPN

An encrypted tunnel between two or more Security Gateways. Synonym: Site-to-Site VPN. Contractions: S2S VPN, S-to-S VPN.

SmartConsole

Check Point GUI application used to manage a Check Point environment - configure Security Policies, configure devices, monitor products and events, install updates, and so on.

SmartDashboard

Legacy Check Point GUI client used to create and manage the security settings in versions R77.30 and lower. In versions R80.X and higher is still used to configure specific legacy settings.

SmartProvisioning

Check Point Software Blade on a Management Server (the actual name is "Provisioning") that manages large-scale deployments of Check Point Security Gateways using configuration profiles. Synonyms: Large-Scale Management, SmartLSM, LSM.

SmartUpdate

Legacy Check Point GUI client used to manage licenses and contracts in a Check Point environment.

Software Blade

Specific security solution (module): (1) On a Security Gateway, each Software Blade inspects specific characteristics of the traffic (2) On a Management Server, each Software Blade enables different management capabilities.

Standalone

Configuration in which the Security Gateway and the Security Management Server products are installed and configured on the same server.

T

Threat Emulation

Check Point Software Blade on a Security Gateway that monitors the behavior of files in a sandbox to determine whether or not they are malicious. Acronym: TE.

Threat Extraction

Check Point Software Blade on a Security Gateway that removes malicious content from files. Acronym: TEX.

U

Updatable Object

Network object that represents an external service, such as Microsoft 365, AWS, Geo locations, and more.

URL Filtering

Check Point Software Blade on a Security Gateway that allows granular control over which web sites can be accessed by a given group of users, computers or networks. Acronym: URLF.

User Directory

Check Point Software Blade on a Management Server that integrates LDAP and other external user management servers with Check Point products and security solutions.

V

VPN Community

A named collection of VPN domains, each protected by a VPN gateway.

VPN Tunnel

An encrypted connection between two hosts using standard protocols (such as L2TP) to encrypt traffic going in and decrypt it coming out, creating an encapsulated network through which data can be safely shared as though on a physical private line.

VSX

Virtual System Extension. Check Point virtual networking solution, hosted on a computer or cluster with virtual abstractions of Check Point Security Gateways and other network devices. These Virtual Devices provide the same functionality as their physical counterparts.

VSX Gateway

Physical server that hosts VSX virtual networks, including all Virtual Devices that provide the functionality of physical network devices. It holds at least one Virtual System, which is called VS0.

Z

Zero Phishing

Check Point Software Blade on a Security Gateway (R81.20 and higher) that provides real-time phishing prevention based on URLs. Acronym: ZPH.

Check Point VPN

IPsec VPN

The IPsec VPN solution lets the Security Gateway encrypt and decrypt traffic to and from other Security Gateways and clients. Use SmartConsole to easily configure VPN connections between Security Gateways and remote devices.

For Site-to-Site Communities, you can configure Star and Mesh topologies for VPN networks, and include third-party gateways.

The VPN tunnel guarantees:

- Authenticity - Uses standard authentication methods
- Privacy - All VPN data is encrypted
- Integrity - Uses industry-standard integrity assurance methods

IKE and IPsec

The Check Point VPN solution uses these secure VPN protocols to manage encryption keys, and send encrypted packets. IKE (Internet Key Exchange) is a standard key management protocol that is used to create the VPN tunnels. IPsec is protocol that supports secure IP communications that are authenticated and encrypted on private or public networks.

VPN Components

VPN is composed of:

- VPN endpoints, such as Security Gateways, Security Gateway clusters, or remote clients (such as laptop computers or mobile phones) that communicate over a VPN.
- VPN trust entities, such as a Check Point Internal Certificate Authority (ICA). The ICA is part of the Check Point suite used for creating SIC trusted connection between Security Gateways, authenticating administrators and third party servers. The ICA provides certificates for internal Security Gateways and remote access clients which negotiate the VPN link.
- VPN Management tools, such as Security Management Server and SmartConsole. The SmartConsole lets organizations define and deploy Intranet, and remote Access VPNs.

Understanding the Terminology

- **VPN** - Virtual Private Network. A secure, encrypted connection between networks and remote clients on a public infrastructure, to give authenticated remote users and sites secured access to an organization's network and resources.
- **Virtual Tunnel Interface** -Virtual Tunnel Interface. A virtual interface that is a member of an existing, Route Based, VPN tunnel.
- **VPN Peer** - A gateway that connects to a different VPN gateway using a Virtual Tunnel Interface.
- **VPN Domain** - A group of computers and networks connected to a VPN tunnel by one VPN Gateway that handles encryption and protects the VPN Domain members.
- **VPN Community** - A named collection of VPN domains, each protected by a VPN Gateway.
- **VPN Security Gateway** - The Security Gateway that manages encryption and decryption of traffic between members of a VPN Domain, typically located at one (Remote Access VPN) or both (Site to Site VPN) ends of a VPN tunnel.
- **Site to Site VPN** - An encrypted tunnel between two Security Gateways, typically of different geographical sites.
- **Remote Access VPN** - An encryption tunnel between a Security Gateway and Remote Access clients, such as Endpoint Security VPN, and communities.
- **Remote Access Community** - A group of computers, appliances, and devices that access, with authentication and encryption, the internal protected network from physically remote sites.
- **Star Topology** - A "hub and spoke" virtual private network community, with Security Gateways defined as Satellites (spokes) that create tunnels only with the central Security Gateway ("hub").
- **Meshed topology** - A VPN community with a VPN Domain that creates a tunnel to other VPN Domains.
- **Domain-based VPN** - A method to route encrypted traffic with parameters defined by Security Gateways.
- **Route-based VPN** - A routing method for participants in a VPN community, defined by the Virtual Tunnel Interfaces (VTI).
- **IKE (Internet Key Exchange)** - An Encryption key management protocol that enhances IPsec by providing additional features, flexibility, and ease of configuration.
- **IPsec** - A set of secure VPN protocols that manage encryption keys and encrypted packet traffic, to create a standard for authentication and encryption services.

Site-to-Site VPN


The basis of Site-to-Site VPN is the encrypted VPN tunnel. Two Security Gateways negotiate a link and create a VPN tunnel and each tunnel can contain more than one VPN connection. One Security Gateway can maintain more than one VPN tunnel at the same time.

Sample Site-to-Site VPN Deployment

Item	Description
A, B	Security Gateways
2	VPN tunnel
3	Internal network in VPN domain
4	Host 4
5	Host 5

In this sample VPN deployment, Host 4 and Host 5 securely send data to each other. The Security Gateways perform IKE negotiation and create a VPN tunnel. They use the IPsec protocol to encrypt and decrypt data that is sent between Host 4 and Host 5.

VPN Workflow




VPN Communities

A VPN Domain is a collection of internal networks that use Security Gateways to send and receive VPN traffic. Define the resources that are included in the VPN Domain for each Security Gateway. Then join the Security Gateways into a VPN community - collection of VPN tunnels and their attributes. Network resources of different VPN Domains can securely communicate with each other through VPN tunnels that terminate at the Security Gateways in the VPN communities.


VPN communities are based on Star and Mesh topologies:

- In a Star community, each satellite Security Gateway has a VPN tunnel to the central Security Gateway, but not to other Security Gateways in the community.
- In a Mesh community, there are VPN tunnels between each pair of Security Gateway.

Mesh Topology




Star Topology



Item	Description
1	Security Gateway
2	Satellite Security Gateways
3	Central Security Gateway

Sample Combination VPN Community



Item	Description
1	London Security Gateway
2	New York Security Gateway
3	London - New York Mesh community
4	London company partner (external network)
5	London Star community
6	New York company partner (external network)
7	New York Star community

This deployment is composed of a Mesh community for London and New York Security Gateways that share internal networks. The Security Gateways for external networks of company partners do not have access to the London and New York internal networks. However, the Star VPN communities let the company partners access the internal networks of the sites that they work with.

Routing VPN Traffic

Configure the Security Gateway to route VPN traffic based on VPN Domains or based on the routing settings of the operating system.

Note - For each VPN Security Gateway, you must configure an existing Security Gateway as a default gateway.

Domain Based VPN

The VPN traffic is routed according to the VPN Domains that are defined in SmartConsole. Use domain based routing to let satellite Security Gateways in a star-based topology send VPN traffic to each other. The central Security Gateway creates a VPN tunnel to each satellite Security Gateway and the traffic is routed to the correct VPN domain.

Route Based VPN

VPN traffic is routed according to the routing settings (static or dynamic) of the Security Gateway operating system. The Security Gateway uses a VTI (VPN Tunnel Interface) to send the VPN traffic as if it were a physical interface. The VTIs of Security Gateways in a VPN community connect and can support dynamic routing protocols.

Granular Routing Control

The Link Selection feature gives you granular control of the VPN traffic in the network. Use this feature to enable the Security Gateway to:

- Find the best possible route for VPN traffic
- Select the interfaces that are used for VPN traffic to internal and external networks
- Configure the IP addresses that are used for VPN traffic
- Use route probing to select available VPN tunnels
- Use Load Sharing for Link Selection to equally distribute VPN traffic to VPN tunnels

IPv6 Support and Limitations

This release includes limited IPv6 support for IPsec VPN communities:

- IPv6 is supported for Site to Site VPN only (Main IP to Main IP).
The Main IP address for both Security Gateways must be defined as an IPv6 Address. You can define other IP addresses that are IPv4 or IPv6.
- IPv6 supports IKEv2 encryption only. IKEv2 is automatically always used for IPv6 traffic.
You can configure the encryption method only for IPv4 traffic.
- VPN tunneling only supports IPv4 inside an IPv4 tunnel, and IPv6 inside an IPv6 tunnel.
IPv4 traffic inside an IPv6 tunnel is not supported.

These VPN features are **not** supported for IPv6:

- Remote Access VPN
- CRL fetch for the Internal Certificate Authority
- Multiple Entry Points (MEP)

- Route-based VPN (VTI)
- Wire Mode VPN
- Security Gateways with a dynamic IP address (DAIP).
- Route Injection Mechanism (RIM)
- Traditional mode Firewall Policies
- IKE Denial of Service protection
- IKE Aggressive Mode
- Security Gateways with Dynamic IP addresses (DAIP)
- Traditional Mode VPN
- Migration from Traditional mode to Simplified mode
- Tunnel Management (permanent tunnels)
- Directional VPN Enforcement
- Link Selection
- GRE Tunnels
- Tunnel View in SmartView Monitor
- VPN Overview page
- `$FWDIR/conf/vpn_route.conf` configuration file

Getting Started with Site-to-Site VPN

Step 1 - Enable the IPsec VPN Software Blade on Security Gateways

Site to Site VPN requires two or more Security Gateways with the IPsec VPN Software Blade enabled. Other Software Blades can be enabled on these Security Gateways.

Make sure that Trusted Communication is established between all Security Gateways and the Management Server.

Do these steps in SmartConsole:

1. Create the Security Gateway objects.
2. Create the Trusted Communication (SIC) with the Management Server.
3. Enable the IPsec VPN Software Blade.

On the **General Properties** page, in the **Network Security** tab, select **IPsec VPN**.

4. Click **OK**.

Note - An internal CA certificate for the Security Gateway is created automatically.

Step 2 - Create a VPN Community

You can create a Meshed or Star VPN Community. See ["VPN Communities" on page 28](#).

The procedure below shows an example of a Star Community.

Configuring a new VPN community

1. From the left navigation panel, click **Security Policies**.
2. In the top left section **Access Control**, click **Policy**.
3. In the bottom left section **Access Tools**, click **VPN Communities**.
4. Click **New** (✱) and select **Star Community**.
5. Enter a name for the VPN Community.
6. In the **Center Gateways** area, click the + icon to add one or more Security Gateways (Clusters) to be in the center of the community.
7. In the **Satellite Gateways** area, click the + icon to add one or more Security Gateways (Clusters) to be around the center Security Gateways (Clusters).
8. Click **OK**.

The Community uses the default encryption and VPN Routing settings.

9. **Optional:** Edit more settings for the VPN Community in the community object.

More VPN Community Settings

In addition to the Security Gateway members, you can edit these settings for the VPN Community in the community object:

- **Encrypted Traffic** - Select **Accept all encrypted traffic** to encrypt and decrypt all traffic between the Security Gateways. If this is not selected, create rules in the Security Policy Rule Base to allow encrypted traffic between community members
- **Encryption** - Select encryption settings that include the **Encryption Method** and **Encryption Suite**. See ["VPN Community Object - Encryption Settings" on page 65](#).
- **Tunnel Management** - Select settings VPN tunnels that include **Permanent Tunnels** and Tunnel Sharing. See ["Configuring Tunnel Features" on page 149](#).
- **VPN Routing** -For Star Communities, select how VPN traffic is routed between the center and satellite Security Gateways. By default this is always set to **To center only**. See ["Configuring VPN Routing in Domain Based VPN" on page 114](#).
- **MEP (Multiple Entry Points)** - For Star Communities, select how the entry Security Gateway for VPN traffic is chosen. This only applies when you have multiple center Security Gateways in the community. See ["Overview of MEP" on page 174](#).
- **Excluded Services** - Add services that are **not** to be encrypted, for example Check Point Control Connections. VPN tunnels are not created for the Services included here.
- **Shared Secret** - Configure shared secret authentication to use for communication with external Security Gateways that are part of a VPN community. See ["Configuring a VPN with External Security Gateways Using Pre-Shared Secret" on page 50](#).
- **Wire Mode** - Select to define internal interfaces and communities as trusted and bypass the Security Gateway for some communication. See ["Configuring Wire Mode" on page 168](#).
- **Advanced** - Configure advanced settings related to IKE, IPsec, and NAT. You can also **Reset All VPN Properties** to revert all VPN Community settings to their default values. See ["Configuring Advanced IKE Properties" on page 65](#).

Step 3 - Configure the VPN Domain for Security Gateways

The VPN Domain defines the networks and IP addresses that are included in the VPN community. It is also called the Encryption Domain. When you create a Check Point Security Gateway object, the VPN Domain is automatically defined as all IP Addresses behind the Security Gateway, based on the topology information.

You can manually define the VPN domain to include one or more networks behind the Security Gateway. You must have a **Network** object or a **Network Group** object that represents the Domain.

Configuring a VPN Domain manually

1. In SmartConsole, from the **Gateways & Servers** view, open a Security Gateway object.
2. Open the **Network Management > VPN Domain** page.
3. Select **Manually defined** and:
 - Browse to the object list and select an object that represents the domain.
 - Browse to the object list and click **New > Group** or **Network** to define a new group of hosts or networks.
4. Click **OK**.

Encryption Domain per Community



Important - This feature requires Security Gateways versions R80.40 and higher.


By default a gateway's Encryption Domain is shared with all the communities it is a part of.

Access to different resources within the Encryption Domain is implemented using the Access Control Rule Base.

In some cases you may need to configure the Encryption Domain in a granular way.

You can configure the VPN domain of a Security Gateway per community, which makes it safer and easier to control the VPN communities that are logically separated.

Example 1



- Security Gateway A (Partner A) is part of Community-1.
- Security Gateway B (Partner B) is part of Community-2.
- Security Gateway C (Corporate Branch) is part of both Communities 1 and 2.
- The network behind Security Gateway C - 10.2.2.0 is split into 2 networks using the 255.255.255.128 subnet mask.

In this scenario, the administrator limits the access from *Security Gateway A* in *community 1* to some of the resources behind *Security Gateway C* which is also part of *community 1*.


To allow access to the required resources from Security Gateway A to resources protected by Security Gateway C, the administrator configures an *Encryption Domain* per the specific community so although Security Gateway C is a part of another community (Community 2) which is configured differently.

The access is limited to the specific *Encryption Domain*: network 10.2.2.0/25.

Configuring the Encryption Domain per Community

1. Open the VPN community.
2. Double click the center Security Gateway that participates in more than one VPN community (Security Gateway C in this scenario).

The VPN domain configuration window opens.



3. Select the **User defined** option.

Configure the Encryption Domain. In our example the encryption domain includes the network we allow partner B to access.

Network

vpn-domain-partner-a
Enter Object Comment

General

NAT

Groups

IPv4

Network address: 10.2.2.0

Net mask: 255.255.255.128

Broadcast address:

☒ Included

☐ Not included

IPv6

Network address:

Prefix:

Add Tag

OK Cancel

4. Click **OK**.

Example 2

Using the same setup, you can use the Encryption Domain per Community configuration to allow access between *host 1* and *host 2* in both directions.

The configuration changes are applied to the Encryption Domain of *Security Gateway-C* per each relevant community, in this example Communities 1 and 2.

Note - In previous versions to get this functionality the `vpn_route.conf` file was used.

Configuration:

Community	Encryption Domain	Install On
Community-1	Network behind Security Gateway-C Host-2	Security Gateway-C Security Gateway-A
Community-2	Network behind Security Gateway-C Host-1	Security Gateway-C Security Gateway-B

1. Create a new host (*Host-2* behind Security Gateway-B) to represent the Encryption Domain of Security Gateway-C to publish for Security Gateway-A.
2. Create a new host (*Host-1* behind Security Gateway-A) to represent the Encryption Domain of Security Gateway-C to publish for Security Gateway-B.
3. Create a new Network group to include the current Encryption Domain of Security Gateway-C and the additional host (*Host-2*) for *Community-1*.
4. Create a new Network group to include the current Encryption Domain of Security Gateway-C and the additional host (*Host-1*) for *Community-2*.
5. For *Community-1* change the Encryption Domain for Security Gateway-C, use the new group created in step 3.
6. For *Community-2* change the Encryption Domain for Security Gateway-C, use the new group created in step 4.

In practice this type of configuration "tricks" the satellite gateways to think that the destination host is part of Security Gateway-C 's Encryption Domain and therefore encrypt the packets from the satellite gateways towards the center Security Gateway. When the encrypted packet gets to the center Security Gateway, it is decrypted and re-routed to its original destination thus it is encrypted again and sent to the other satellite gateway.

Specific VPN Domain for Gateway Communities

If a Security Gateway participates in more than one VPN Community, you can configure a different VPN Domain for the Security Gateway for each VPN Community in which it participates. In SmartConsole, you can configure a specific VPN Domain for a Security Gateway in the Security Gateway object or in the VPN Community object.

To configure a specific VPN Domain in the Security Gateway Object:

1. Open the **Network Management > VPN Domain** page.
2. In the line **Set Specific Domain for Gateway Communities**, click **Set**.
3. Select the VPN Community for which it is necessary to override the VPN Domain and click **Set**.
4. Select the applicable option:
 - **According to the gateway**

This configuration option use the VPN Domain that is configured in the **Network Management** folder > **VPN Domain** page > **VPN Domain** section.

■ User defined

Select the applicable Network or Group object (or create a new object).

This configuration option overrides:

- The VPN Domain that is configured in the Security Gateway object > **Network Management** folder > **VPN Domain** page > **VPN Domain** section.
- The VPN Domain that is configured in the Meshed / Star VPN Community object > **Gateways** page.
- The VPN Domain that is configured in the Remote Access VPN Community object > **Participating Gateways** page.

5. Click **OK** to close the Set Specific VPN Domain for Gateway Communities window.
6. Click **OK** to close the **Communities Specific VPN Domain** window.

To configure a specific VPN Domain in the VPN Community Object:

1. In the **Objects** pane, click **VPN Communities**.
2. Click the applicable VPN Community.
The VPN Community configuration window opens.
3. In the **Gateways** pane, double-click the relevant Security Gateway object (or create a new object).
The VPN Domain configuration window opens.
4. Select the applicable option:

■ According to the gateway

This configuration option use the VPN Domain that is configured in the **Network Management** folder > **VPN Domain** page > **VPN Domain** section.

■ User defined


Select the applicable Network or Group object (or create a new object).

This configuration option overrides:

- The VPN Domain that is configured in the Security Gateway object > **Network Management** folder > **VPN Domain** page > **VPN Domain** section.
 - The VPN Domain that is configured in the Meshed / Star VPN Community object > **Gateways** page.
 - The VPN Domain that is configured in the Remote Access VPN Community object > **Participating Gateways** page.
5. Click **OK** to close the VPN Domain configuration window.
 6. Click **OK** to close the VPN Community configuration window.

Granular Encryption for Externally Managed Gateways

The need for Granular Encryption - Many times organizations are required to connect a third party VPN Gateway to an existing VPN community, and for security reasons requires the use of a stronger encryption suite. With Granular Encryption you can add an Externally Managed Gateway that uses a different encryption suite to participate in an existing community without the need to change the encryption methods in use or split the VPN community.



Note - Granular Encryption can be used only with Security Gateways that run R81 or higher.


Overriding Default Encryption Methods

1. Open the VPN community.
2. Select the **Encryption settings** tab.

- At the bottom of the settings window, beneath the **Override Encryption for Externally Managed Gateways** click the **+** button.

The **Granular Encryption** window opens.

Example:



- Select the Security Gateways that connects with the Externally Managed Gateway.

Granular Encryption settings are set in pairs, the Internal Security Gateway and the Externally Managed Security Gateway that corresponds, this is the **Encryption Context**.

The default value for the **Internal Gateway** is *** Any**. If this option is used, all the Internal Gateways participating in the VPN community use the same Encryption Suite to establish the VPN connection with the Externally Managed Gateway.

Note - If Granular Encryption is set for a specific Internal Gateway in addition to the use of *** Any** in a different Encryption Context, the Granular Encryption settings apply.

- Select the **Encryption Method** and **Encryption Suite** to use for the VPN communication between the selected peers.

Step 4 - Make Sure VPN Routing Works

Make sure the VPN works with the routing configured in your network. If it does not work, change the routing configuring or change the **Link Selection** settings as necessary. See ["Link Selection Overview" on page 69](#).

By default, IPsec VPN uses the main **IPv4 Address**, defined in the **General Properties** page of the Security Gateway object, for the VPN tunnel connection.

If you want to use this IP address for the VPN communication, and it is an external interface, you do not need additional routing.

If the main IP address is an internal interface, or if you want VPN communication on a different interface, make sure that:

- The **Link Selection** settings for the Security Gateway are configured. Choose which Security Gateway links are used by VPN to route traffic correctly. See ["Link Selection Overview" on page 69](#)
- VPN Routing is configured to allow the connections. For information how to configure routing in Gaia OS, see the [R81.10 Gaia Administration Guide](#) - Chapter *Network Management*.

Step 5 - Configure the Access Control Rules and install the Access Control Policy

You must configure Access Control rules to allow traffic within VPN Communities.

Configure rules in SmartConsole > **Security Policies** view > **Access Control**.

All layers of the Access Control Policy can contain VPN rules.

To make a rule apply to a VPN Community, the **VPN** column of the Rule Base must contain one of these:

- **Any** - The rules applies to all VPN Communities and to non-VPN related traffic. If you configure a new VPN Community after the rule was created, the rule also applies to the new VPN Community.
- **One or more specified VPN communities** - For example, **MyIntranet**. Right-click in the VPN column of a rule and select **Specific VPN Communities**. The rule applies to the communities shown in the VPN column.

Examples:

- This rule allows encrypted traffic between domains of member Security Gateways of "community_X."

Name	Source	Destination	VPN	Services & Applications
Allow traffic within community	* Any	*Any	✱ MyCommunity	* Any

- This rule allows traffic from all VPN Communities to the internal network on all services.

Name	Source	Destination	VPN	Services & Applications
Allow all VPN	* Any	Internal_Network	* Any	* Any

- This rule allows traffic between two VPN domains with all services.

Name	Source	Destination	VPN	Services & Applications
Site to Site VPN	Local_VPN_Domain Peer_VPN_Domain	Local_VPN_Domain Peer_VPN_Domain	✱ Site2Site	* Any


Step 6 - Test the VPN Tunnel

To make sure that a VPN tunnel works:

1. Locate the Access Control rule for the traffic that has to pass through the VPN tunnel.
In the **Track** column, select **Log**.
2. From the left navigation panel, click **Logs & Monitor > Logs**.
3. From the top, click **New Tab**.
4. From the bottom of the window, click **Tunnel and User Monitoring**.
Check Point SmartView Monitor opens.
5. Click the Security Gateway to see IPsec VPN traffic and tunnels opened.
A successful connection shows encrypt, decrypt and key install logs.
Alternatively:

- a. In SmartConsole, from the left navigation panel, click **Logs & Monitor**.
- b. On the **Logs** tab, search for **VPN** to see the applicable logs.

Example:



The screenshot shows the SmartConsole interface with the 'Logs' tab selected. The interface includes a search bar with a dropdown menu set to 'Last 7 Days' and a search query input field. Below the search bar, it indicates 'Showing first 100 results (962 ms) out of 2,838 results'. The main area displays a table of log entries.

Time	B...	I...	Origin	A...	Source	Destination	Service	Policy...	Description
Today, 7:46:32 AM			GW124		GW124 (172.23.43.124)	GW126 (172.23.43.126)			
Today, 7:46:32 AM			GW126		GW124 (10.2.2.1)	GW126 (10.2.2.10)	echo-request	Standard	Decrypted in community MyIntranet
Today, 7:46:32 AM			GW124		GW124 (172.23.43.124)	GW126 (172.23.43.126)			
Today, 7:46:32 AM			GW124		GW124 (172.23.43.124)	GW126 (172.23.43.126)	IKE	Standard	IKE Traffic Accepted from 172.23.43.124 to 172.23.
Today, 7:46:32 AM			GW126		GW124 (172.23.43.124)	GW126 (172.23.43.126)			
Today, 7:46:32 AM			GW124		GW124 (10.2.2.1)	GW126 (10.2.2.10)	echo-request	Standard	Encrypted in community MyIntranet
Today, 7:46:32 AM			GW126		GW124 (172.23.43.124)	GW126 (172.23.43.126)			
Today, 7:46:32 AM			GW126		GW124 (172.23.43.124)	GW126 (172.23.43.126)	IKE	Standard	IKE Traffic Accepted from 172.23.43.124 to 172.23.

Basic Site to Site VPN Configuration

It is more complex to configure VPN with external Security Gateways (those managed by a different Security Management Server) than to configure VPN with internal Security Gateways (managed by the same Security Management Server) because:

- There are two systems to configure separately.
- Administrators of the peer VPN Security Gateways must coordinate with each other and agree on all details. The administrators must manually supply details such as the IP address and the VPN domain topology. These details cannot be detected automatically.

Configuring a Star or Meshed Community Between Internally Managed Security Gateways

See ["VPN Communities" on page 28](#).

Procedure

1. Install and configure the Security Gateways as described in the [R81.10 Installation and Upgrade Guide](#).
 2. From the left navigation panel, click **Gateways & Servers**.
 3. Open the Security Gateway object.
 4. On the **General Properties** page, click the **Network Security** tab, and select **IPsec VPN**.
 5. Configure the VPN Domain:
 - a. From the left tree, click **Network Management > VPN Domain**.
 - b. Select one of these:
 - **All IP Addresses behind the Gateway based on Topology information**
 - **User-defined** - select the applicable object (Network, Address Range, Group).
-  **Note** - There is nothing to configure on the **IPsec VPN** page for certificates. This is because Security Gateways that this Management Server manages automatically receive a certificate from this Management Server's Internal Certificate Authority.
6. Click **OK**.

7. From the top toolbar, click **Objects > Object Explorer**.
8. From the left tree, click **VPN Communities**.
9. Create a new VPN Community object.

To create a Star community

Click **New > VPN Community > Star Community**.

The **New Star Community** window opens.

- a. Enter the name for this VPN Community.
- b. On the **Gateways** page:
 - In the **Center Gateways** section, select the applicable Security Gateway / Cluster objects.



Important - This field does not support:

- VSX Gateways and VSX Clusters.
- Maestro Security Groups.
- Quantum Spark appliances that run Gaia Embedded OS.

Select **Mesh center gateways** for the center Security Gateways to connect with each other.

- In the **Satellite Gateways** section, select the applicable Security Gateway objects.

- c. On the **Encrypted Traffic** page:

Select **Accept all encrypted traffic**, if it is necessary to encrypt all traffic between the Security Gateways.

Select the applicable option:

- **Both center and satellite gateways**
- **Satellite gateways only**

If you do not need to encrypt all traffic between the Security Gateways, then create the applicable Access Control rules in the Security Policy (see the next step).

- d. On the **VPN Routing** page, select **To center only**.
- e. Click **OK**.
- f. Close the Object Explorer window.

For information on other options, such as **Encryption**, **Shared Secret**, and **Advanced**, see ["IPsec and IKE" on page 55](#).

For information on the **MEP** option, see ["Multiple Entry Point \(MEP\) VPNs" on page 174](#).

To create a Meshed community

Click **New > VPN Community > Meshed Community**.

The **New Meshed Community** window opens.

a. Enter the name for this VPN Community.

b. On the **Gateways** page:

Add the applicable Security Gateway objects.

c. On the **Encrypted Traffic** page:

Select **Accept all encrypted traffic**, if it is necessary to encrypt all traffic between the Security Gateways.

If you do not need to encrypt all traffic between the Security Gateways, then create the applicable Access Control rules in the Security Policy (see the next step).

d. Click **OK**.

e. Close the Object Explorer window.

For information on other options, such as **Encryption**, **Shared Secret**, and **Advanced**, see ["IPsec and IKE" on page 55](#).

10. If you did not select **Accept all encrypted traffic** on the **Encrypted Traffic** page of the VPN Community, configure the applicable Access Control rules.

For example:

Source	Destination	VPN	Service	Action
* Any	* Any	Meshed VPN Community you configured	* Any	Accept

For more information on how to configure an Access Control policy, see the [R81.10 Security Management Administration Guide](#).

11. Install the Access Control Policy on these Security Gateways.

Configuring a VPN with External Security Gateways Using Certificates

This section applies to typical configurations of a VPN with External Security Gateways, and assumes that the peers work with certificates. If this is not the case, see ["Configuring a VPN with External Security Gateways Using Pre-Shared Secret" on page 50](#).

To configure a VPN with an externally managed peer, you and the peer administrator must choose the same Certificate Authority (CA) for communication between the two peers.

Even if each of the peer VPN Security Gateways uses a Check Point Internal CA (ICA), if they are not managed by the same Security Management Server then their ICAs are different.

Example - A Check Point Security Gateway located at a headquarters office and a peer Check Point Security Gateway located at a branch office are managed separately. Each peer Security Gateway uses a different Check Point ICA and has different parameters for encryption. The administrators of the two networks must agree on a CA for communication between the two peers.

Note - Configuring a VPN with PKI and certificates is more secure than with pre-shared secrets.

Procedure

1. Get the certificate of the CA that issued the certificate for the peer VPN Security Gateways. Request this from the peer administrator.

If the peer Security Gateway uses the Internal Certificate Authority, then to obtain the Certificate Authority certificate file, connect with a web browser to this portal:

- In R81.10 and higher:

```
http://<IP address of Management Server that manages  
the peer Security Gateway>:18268
```

- In R81 and lower:

```
http://<IP address of Management Server that manages  
the peer Security Gateway>:18265
```

2. In SmartConsole, configure the Certificate Authority object for the Certificate Authority that issued the certificate for the peer. See ["Enrolling with a Certificate Authority" on page 99](#).
3. Configure a Certificate Authority to issue certificates for your side in case the Certificate issued by ICA is not applicable for the required VPN tunnel.

You may have to export the CA certificate and supply it to the peer administrator.

4. Define the Network Object(s) of the Security Gateway(s) that are internally managed:
 - In the **General Properties** page of the Security Gateway object, select **IPsec VPN**.
 - In the **Network Management** page, define the **Topology**.
 - In the **VPN Domain** page, define the **VPN Domain**.
 If the VPN domain does not contain all the IP addresses behind the Security Gateway, then configure the VPN domain manually by defining a group or network of machines and setting them as the VPN Domain.
5. If the ICA certificate is not applicable for this VPN tunnel, then generate a certificate from the applicable Certificate Authority on the **IPsec VPN** page.
6. Define the Network Object(s) of the externally managed Security Gateway(s).
 - If it is not a Check Point Security Gateway, define an Interoperable Device:
 In **Object Explorer**, click **New > Network Object > More > Interoperable Device**.
 - If it is a Check Point Security Gateway, define an Externally Managed VPN Gateway:
 In **Object Explorer**, click **New > Network Object > Gateways and Servers > More > Externally Managed VPN Gateway**.
7. Set the attributes of the peer Security Gateway.
 - For an externally managed Check Point Security Gateway:
 In the **General Properties** page of the Security Gateway object, select **IPsec VPN**.
 - Define the **Topology**.
 - Define the **VPN Domain** with the VPN Domain information obtained from the peer administrator. If the VPN Domain does not contain all the IP addresses behind the Security Gateway, define the VPN domain manually by defining a group or network of machines and setting them as the VPN Domain.
 - For an Externally Managed Check Point Security Gateway:
 On the IPsec VPN page, define the **Matching Criteria**. Specify that the peer must present a certificate signed by its own Certificate Authority. If possible, enforce details that appear in the certificate.
8. Define the VPN Community.

If you are configuring a meshed community rather than a star community, ignore the difference between the Central Security Gateways and the Satellite Security Gateways.

- Agree with the peer administrator about the various IKE properties and set them in the **Encryption** page and the **Advanced** page of the community object.
 - Define the Central Security Gateways. In most cases these are internal. If no other Community is defined for them, decide whether to mesh the central Security Gateways. If they are already in a Community, do not mesh the Central Security Gateways.
 - Define the Satellite Security Gateways. In most cases these are external.
9. Click **OK**.
 10. Publish the SmartConsole session.
 11. Define the applicable Access Control rules.
 12. Add the Community in the **VPN** column, the services in the **Service & Applications** column, the **Action**, and the applicable **Track option**.
 13. Install the Access Control Policy.

Configuring a VPN with External Security Gateways Using Pre-Shared Secret

Administrators of the peer VPN Security Gateways must coordinate with each other and agree on all details. The administrators must manually supply details such as the IP address and the VPN domain topology. These details cannot be detected automatically.

There are many possible scenarios for VPN with external Security Gateways. The next procedure is meant for typical cases and assumes that the peers work with pre-shared secrets. If the peers do not work with pre-shared secrets, see **Configuring a VPN with External Security Gateways Using Certificates**.

 **Note** - It is more secure to configure a VPN with public key infrastructure (PKI) and certificates than with pre-shared secrets.

To configure a VPN using pre-shared secrets with the external Security Gateways as satellites in a star VPN Community:

1. Define the Network Object(s) of the Security Gateways that are internally managed.
 - In the **General Properties** page of the Security Gateway object, in the Network Security tab, select IPsec VPN.
 - In the **Network Management** page, define the Topology.

- In the **Network Management > VPN Domain** page, define the VPN Domain. If the VPN domain does not contain all IP addresses behind the Security Gateway, define the VPN Domain manually by defining a group or network of machines and setting them as the VPN Domain.
- 2. Define the Network Object(s) of the externally managed Security Gateway(s).
 - If it is not a Check Point Security Gateway, define an Interoperable Device:
In Object Explorer, click **New > Network Object > More > Interoperable Device**.
 - If it is a Check Point Security Gateway, define an Externally Managed VPN Gateway:
In Object Explorer, click **New > Network Object > Gateways and Servers > More > Externally Managed VPN Gateway**.
- 3. Set the attributes of the peer Security Gateway.
 - In the **Topology** page, define the Topology and the VPN Domain with the VPN Domain information obtained from the peer administrator.
 - If the VPN Domain does not contain all the IP addresses behind the Security Gateway, configure the VPN Domain manually by defining a group or network of machines and setting them as the VPN Domain.
- 4. Define the Community.

If you are configuring a Mesh Community rather than a Star Community, ignore the difference between the Central Security Gateways and the Satellite Security Gateways.

 - Agree with the peer administrator about the IKE properties. Set the IKE properties in the **Encryption** page and the **Advanced** page of the community object.
 - Define the Central Security Gateways. These are usually the internally managed Security Gateways. If there is not another Community defined for them, decide whether to mesh the central Security Gateways. If the Central Security Gateways are already in a Community, do not mesh them.
 - Define the Satellite Security Gateways. These are usually the external Security Gateways.
- 5. Publish the changes in SmartConsole.

6. Agree on a pre-shared secret with the administrator of the external Community members. Then, in the **Shared Secret** page of the Community, select **Use only Shared Secret** for all external members. For each external member, enter the pre-shared secret.
7. Define the applicable Access Control rules in the **Access Control Policy**. Add the Community in the **VPN** column, the services in the **Services & Applications** column, the desired **Action**, and the applicable **Track** option.
8. Install the Access Control Policy.

Firewall Control Connections in VPN Communities

Check Point Nodes communicate with other Check Point Nodes through control connections. For example a Security Management Server and a Security Gateway use a control connection when the Security Policy is installed from the Security Management Server to the Security Gateway. In addition, Security Gateways send logs to the Security Management Server across control connections. Control connections use Secure Internal Communication (SIC).

Implied Rules in the Access Control Rule Base allow the Control connections. The Management Server adds and removes the Implied Rules in the Access Control Rule Base when you select or clear options in the SmartConsole > **Menu** > **Global properties** > **Firewall** page.

Some administrators do not rely on implied rules, and instead define explicit rules in the Access Control Rule Base. Check Point does not support replacing implied rules with explicit rules. See [sk43401](#).

Why Turning off Implied Rules Blocks Firewall Control Connections

If you turn off implicit rules, you may not be able to install an Access Control Policy on a remote Security Gateway. Even if you configure explicit rules rather than implied rules, you may still not be able to install the policy:



To configure a VPN between Security Gateways A and B through SmartConsole, the administrator must install a Policy from the Security Management Server to the Security Gateways.

1. The Security Management Server successfully installs the Policy on Security Gateway A. Security Gateway A recognizes that Security Gateways A and B now belong to the same VPN Community. However, Security Gateway B does not yet have the Policy.
2. The Security Management Server opens a connection to Security Gateway B to install the Policy.
3. Security Gateway A allows the connection because of the explicit rules that allow the control connections. Security Gateway A starts IKE negotiation with Security Gateway B to build a VPN tunnel for the control connection.
4. Security Gateway B cannot negotiate with Security Gateway A because it does not yet have the Policy. Therefore, Policy installation on Security Gateway B fails.

Make sure that control connections do not have to pass through a VPN tunnel.

Allowing Firewall Control Connections Inside a VPN

If you turn off implied rules, make sure that control connections are not changed by the Security Gateways. Add the services that are used for control connections to the Excluded Services page of the Community object. See [sk42815](#) for details.

Note - Although control connections between the Security Management Server and the Security Gateway are not encrypted by the community, they are still encrypted and authenticated with Secure Internal Communication (SIC).

Discovering Which Services are Used for Control Connections

1. In SmartConsole, click **Menu > Global properties**.
2. On the **Firewall** page, select **Control Connections**.
3. Click **OK**.
4. In SmartConsole, from the left panel, click **Security Policies**.
5. Select the applicable Access Control Policy.
6. From the toolbar above the policy, select **Actions > Implied Rules**.

The Implied Policy window opens.

7. Examine the Access Control Rule Base to see what Implied Rules are visible. Note the services used in the Implied Rules.

Simplified and Traditional Modes

By default, VPN configuration works with Simplified mode. Simplified mode uses VPN Communities for Site to Site VPN configuration, as described in this Administration Guide.

Traditional mode is a different, legacy way to configure Site to Site VPN where one of the actions available in the Security Policy Rule Base is **Encrypt**. When **Encrypt** is selected, all traffic between the Security Gateways is encrypted. For details about Traditional Mode, see the [R77 versions VPN Administration Guide](#).

In a policy package, all layers must use the same VPN mode.

IPsec and IKE

Overview

In symmetric cryptographic systems, both communicating parties use the same key for encryption and decryption. The material used to build these keys must be exchanged in a secure fashion. Information can be securely exchanged only if the key belongs exclusively to the communicating parties.

The goal of the *Internet Key Exchange* (IKE) is for both sides to independently produce the same symmetrical key. This key then encrypts and decrypts the regular IP packets used in the bulk transfer of data between VPN peers. IKE builds the VPN tunnel by authenticating both sides and reaching an agreement on methods of encryption and integrity. The outcome of an IKE negotiation is a *Security Association* (SA).

This agreement upon keys and methods of encryption must also be performed securely. For this reason, IKE is composed of two phases. The first phase lays the foundations for the second.

Diffie-Hellman (DH) is that part of the IKE protocol used for exchanging the material from which the symmetrical keys are built. The Diffie-Hellman algorithm builds an encryption key known as a "shared secret" from the private key of one party and the public key of the other. Since the IPsec symmetrical keys are derived from this DH key shared between the peers, at no point are symmetric keys actually exchanged.

For more on Cryptographic Suites for IPsec, see: [RFC 4308](#).


IKE Phase I

During IKE Phase I:

- The peers authenticate, either by certificates or via a pre-shared secret. (More authentication methods are available when one of the peers is a remote access client.)
- A Diffie-Hellman key is created. The nature of the Diffie-Hellman protocol means that both sides can independently create the shared secret, a key which is known only to the peers.
- Key material (random bits and other mathematical data) as well as an agreement on methods for IKE phase II are exchanged between the peers.

In terms of performance, the generation of the Diffie-Hellman Key is slow and heavy. The outcome of this phase is the IKE SA, an agreement on keys and methods for IKE phase II. Figure below illustrates the process that takes place during IKE phase I.

Note - The exact negotiation stages differ between IKEv1 and IKEv2.



IKE Phase II (Quick mode or IPSec Phase)

IKE phase II is encrypted according to the keys and methods agreed upon in IKE phase I. The key material exchanged during IKE phase II is used for building the IPsec keys. The outcome of phase II is the IPsec Security Association. The IPsec SA is an agreement on keys and methods for IPsec, thus IPsec takes place according to the keys and methods agreed upon in IKE phase II.

After the IPsec keys are created, bulk data transfer takes place:

IKEv1 and IKEv2

IKEv2 is supported inside VPN communities working in **Simplified** mode.

IKEv2 is configured in the **VPN Community Properties window > Encryption**. The default setting is **IKEv1 only**. IKEv2 is automatically always used for IPv6 traffic. The encryption method configuration applies to IPv4 traffic only.

To configure IKE settings for Remote Access VPN users in SmartConsole, click **Menu > Global properties > Remote Access > VPN - Authentication and Encryption**.

Note - IKE2 is not supported for Remote Access.-

Methods of Encryption and Integrity

Two parameters are decided during the negotiation:

- Encryption algorithm
- Hash algorithm

Parameter	IKE Phase 1 (IKE SA)	IKE PHASE 2 (IPSec SA)
Encryption	• AES-128 • AES-256 (default) • 3DES • DES • CAST (IKEv1 only)	• AES-128 (default) • AES-256 • 3DES • DES • DES-40CP (IKEv1 only) • CAST (IKEv1 only) • CAST-40 (IKEv1 only) • NULL • AES-GCM-128 • AES-GCM-256
Integrity	• MD5 • SHA1 (default) • SHA-256 • SHA-512 • AES-XCBC • SHA -384	• MD5 • SHA1 (default) • SHA-256 • SHA-512 • AES-XCBC • SHA -384

NULL means perform an integrity check only; *packets are not encrypted*.

Diffie Hellman Groups

The Diffie-Hellman key computation (also known as exponential key agreement) is based on the Diffie Hellman (DH) mathematical groups. A Security Gateway supports these DH groups during the two phases of IKE.

Parameter	IKE Phase 1 (IKE SA)	IKE Phase 2 (IPSec SA)
Diffie Hellman Groups	▪ Group2 (1024 bits) (default) ▪ Group1 (768 bits) ▪ Group5 (1536 bits) ▪ Group14 (2048 bits) ▪ Group19 (256-bit ECP) ▪ Group20 (384-bit ECP)	▪ Group2 (1024 bits) (default) ▪ Group1 (768 bits) ▪ Group5 (1536 bits) ▪ Group14 (2048 bits) ▪ Group19 (256-bit ECP) ▪ Group20 (384-bit ECP)

A group with more bits ensures a key that is harder to break, but carries a heavy cost in terms of performance, since the computation requires more CPU cycles.

Phase I modes

Between Security Gateways, there are two modes for IKE phase I. These modes only apply to IKEv1:

- **Main Mode**
- **Aggressive Mode**

If aggressive mode is **not** selected, the Security Gateway defaults to main mode, performing the IKE negotiation with six packets; aggressive mode performs the IKE negotiation with three packets.

Main Mode is preferred because:

- Main mode is partially encrypted, from the point at which the shared DH key is known to both peers.
- Main mode is less susceptible to Denial of Service (DoS) attacks. In main mode, the DH computation is performed *after* authentication. In aggressive mode, the DH computation is performed parallel to authentication. A peer that is not yet authenticated can force processor intensive Diffie-Hellman computations on the other peer.
- **Note** - Use aggressive mode when a Check Point Security Gateway needs to negotiate with third party VPN solutions that do not support main mode.

When dealing with remote access, IKE has additional modes:

- *Hybrid Mode* that provides an alternative to IKE phase I, where the Security Gateway is allowed to authenticate with certificates and the client via some other means, such as SecurID. For more information on Hybrid mode, see the [R81.10 Remote Access VPN Administration Guide](#).
- *Office Mode* that is an extension to the IKE protocol. Office Mode is used to resolve routing issues between remote access clients and the VPN domain. During the IKE negotiation, a special mode called *config mode* is inserted between phases I and II. During config mode, the remote access client requests an IP address from the Security Gateway. After the Security Gateway assigns the IP address, the client creates a virtual adapter in the Operating System. The virtual adapter uses the assigned IP address. For more information, see the [R81.10 Remote Access VPN Administration Guide](#).

Renegotiating IKE & IPsec Lifetimes

IKE phase I is more processor intensive than IKE phase II, because the Diffie-Hellman keys have to be produced, and the peers authenticated, each time. For this reason, IKE phase I is performed less frequently. However, the IKE SA is only valid for a certain period, after which the IKE SA must be renegotiated. The IPsec SA is valid for an even shorter period, meaning many IKE phase II negotiations take place.

The period between each renegotiation is known as the **lifetime**. Generally, the shorter the lifetime, the more secure the IPsec tunnel (at the cost of more processor intensive IKE negotiations). With longer lifetimes, future VPN connections can be set up more quickly. By default, IKE phase I occurs once a day; IKE phase II occurs every hour but the time-out for each phase is configurable.

Configure the frequency of IKE and IPsec Security Associations in SmartConsole > **Objects** menu > **Object Explorer** > **VPN Communities** > VPN Community object > **Advanced**.

Perfect Forward Secrecy

The keys created by peers during IKE phase II and used for IPsec are based on a sequence of random binary digits exchanged between peers, and on the DH key computed during IKE phase I.

The DH key is computed once, then used a number of times during IKE phase II. Since the keys used during IKE phase II are based on the DH key computed during IKE phase I, there exists a mathematical relationship between them. For this reason, the use of a single DH key may weaken the strength of subsequent keys. If one key is compromised, subsequent keys can be compromised with less effort.

In cryptography, **Perfect Forward Secrecy** (PFS) refers to the condition in which the compromise of a current session key or long-term private key does **not** cause the compromise of earlier or subsequent keys. Security Gateways meet this requirement with a PFS mode. When PFS is enabled, a fresh DH key is generated during IKE phase II, and renewed for each key exchange.

However, because a new DH key is generated during each IKE phase I, no dependency exists between these keys and those produced in subsequent IKE Phase I negotiations. Enable PFS in IKE phase II only in situations where extreme security is required.

The supported DH groups for PFS are: 1, 2, 5, 14, 19, and 20. The default is group 2 (1042 bits).

Configure this in **VPN Community Properties** > **Encryption** > **IKE Security Association (Phase 2)** > **Use Perfect Forward Secrecy**.



Notes:

- The Perfect Forward Secrecy (PFS) feature supports only IPsec and only for Endpoint VPN clients. When the PFS is enabled on a Security Gateway, all non-supported Remote Access VPN clients fail to connect with the error "The user is not defined properly".
- The Perfect Forward Secrecy (PFS) feature uses the same Diffie-Helman (DH) group in Phase 2 as configured for Phase 1 (**SmartConsole** > **Menu** > **Global properties** > **Remote Access** > **VPN - Authentication and Encryption** > **Encryption algorithms** > **Edit** > **Phase 1** > **Use Diffie-Helman group**).

IP Compression

IP compression is a process that reduces the size of the data portion of the TCP/IP packet. Such a reduction can cause significant improvement in performance. IPsec supports the *Flate/Deflate* IP compression algorithm. Deflate is a smart algorithm that adapts the way it compresses data to the actual data itself. Whether to use IP compression is decided during IKE phase II. IP compression is not enabled by default.

IP compression is important for Remote Access client users with slow links.


Security Gateway encryption makes TCP/IP packets appear "mixed up". This kind of data cannot be compressed and bandwidth is lost as a result. If IP compression is enabled, packets are compressed *before* encryption. This has the effect of recovering the lost bandwidth.

Subnets and Security Associations

By default, a VPN tunnel is created for the complete subnets that host computers reside on, and not just for the host computers involved in the communication.


Unique SA Per Pair of Peers

If you disable the **Support Key exchange for subnets** option on each Security Gateway, you can create a *unique* Security Association for a pair of peers.



If the Security Gateway is configured to **Support key exchange for subnets**, but the option is unsupported on the remote peer, when Host A communicates with Host C, a Security Association (**SA 1**) will be negotiated between Host A's subnet and Host C's IP address. The same SA is then used between any host on the 10.10.11.x subnet and Host C.

When Host A communicates with Host B, a separate Security Association (**SA 2**) is negotiated between Host A's subnet and Host B. As before, the same SA is then used between any host in 10.10.11.x subnet and Host B.



When **Support Key exchange for subnets** is not enabled on communicating Security Gateways, then a security association is negotiated between individual IP addresses; in effect, a unique SA per host.

IKE DoS Protection

Understanding DoS Attacks

Denial of Service (DoS) attacks are intended to reduce performance, block legitimate users from using a service, or even bring down a service. They are not direct security threats in the sense that no confidential data is exposed, and no user gains unauthorized privileges. However, they consume computer resources such as memory or CPU.

Generally, there are two kinds of DoS attack. One kind consists of sending malformed (garbage) packets in the hope of exploiting a bug and causing the service to fail. In the other kind of DoS attack, an attacker attempts to exploit a vulnerability of the service or protocol by sending well-formed packets. IKE DoS attack protection deals with the second kind of attack.

IKE DoS Attacks

The IKE protocol requires that the receiving Security Gateway allocates memory for the first IKE Phase 1 request packet that it receives. The Security Gateway replies, and receives another packet, which it then processes using the information gathered from the first packet.

An attacker can send many IKE first packets, while forging a different source IP address for each. The receiving Security Gateway is obliged to reply to each, and assign memory for each. This can consume all CPU resources, thereby preventing connections from legitimate users.

The attacker sending IKE packets can pretend to be a machine that is allowed to initiate IKE negotiations, such as a Check Point Security Gateway. This is known as an identified source. The attacker can also pretend to have an IP address that the receiving Security Gateway does not know about, such as a Remote Access client, or a Check Point Security Gateway with a dynamic IP address. This is known as an unidentified source.

Defense Against IKE DoS Attacks

When the number of simultaneous IKE negotiations handled exceeds the accepted threshold, it concludes that it is either under load or experiencing a Denial of Service attack. In such a case, the Security Gateway can filter out peers that are the probable source of a potential Denial of Service attack. The following sections describe different types of defenses against IKE DoS attacks.

IKE DoS protection is not supported for IPv6 addresses.

SmartConsole IKE DoS Attack Protection Settings

To protect against IKE DoS attacks:

1. In SmartConsole, click **Menu > Global properties > VPN > Advanced**.
2. In the **IKE Denial of Service protection** section, configure these settings:
 - **Support IKE DoS protection from identified source** - The default setting for identified sources is **Stateless**. If the Security Gateway is under load, this setting requires the peer to respond to an IKE notification in a way that proves that the IP address of the peer is not spoofed. If the peer cannot prove this, the Security Gateway does not begin the IKE negotiation.

If the source is identified, protecting using **Puzzles** is over cautious, and may affect performance. A third possible setting is **None**, which means no DoS protection.

- **Support IKE DoS protection from unidentified source** - The default setting for unidentified sources is **Puzzles**. If the Security Gateway is under load, this setting requires the peer to solve a mathematical puzzle. Solving this puzzle consumes peer CPU resources in a way that makes it difficult to initiate multiple IKE negotiations simultaneously.

For unidentified sources, **Stateless** protection may not be sufficient because an attacker may well control all the IP addresses from which the IKE requests appear to be sent. A third possible setting is **None**, which means no DoS protection.

3. Click **OK**.
4. Install the Access Control Policy.

Note - IKE DoS protection is not supported for IPv6 addresses.

Advanced IKE DoS Attack Protection Settings

You can configure the advanced IKE DoS attack protection on the Management Server with the Database Tool (GuiDBEdit Tool) (see [sk13009](#)).

Note - IKE DoS protection is not supported for IPv6.

Parameter	Description	Accepted Values	Default Value
<code>ike_dos_threshold</code>	Determines the percentage of maximum concurrent ongoing negotiations, above which the Security Gateway will request DoS protection. If the threshold is set to 0, the Security Gateway always requests DoS protection.	0 - 100	70
<code>ike_dos_puzzle_level_identified_initiator</code>	Determines the level of the puzzles sent to known peer Security Gateways. This parameter also determines the maximum puzzle level a Security Gateway is willing to solve.	0 - 32	19
<code>ike_dos_puzzle_level_unidentified_initiator</code>	Determines the level of the puzzles sent to unknown peers (such as Remote Access clients and DAIP Security Gateways). This parameter also determines the maximum puzzle level that DAIP Security Gateways and Remote Access clients are willing to solve.	0 - 32	19
<code>ike_dos_max_puzzle_time_gw</code>	Determines the maximum time in milliseconds a Security Gateway is willing to spend solving a DoS protection puzzle.	0 - 30000	500
<code>ike_dos_max_puzzle_time_daip</code>	Determines the maximum time in milliseconds a DAIP Security Gateway is willing to spend solving a DoS protection puzzle.	0 - 30000	500

Parameter	Description	Accepted Values	Default Value
<code>ike_dos_max_puzzle_time_sr</code>	Determines the maximum time in milliseconds a client is willing to spend solving a DoS protection puzzle.	0 - 30000	5000
<code>ike_dos_supported_protection_sr</code>	When downloaded to a client, it controls the level of protection the client is willing to support. Security Gateways use the <code>ike_dos_protection_unidentified_initiator</code> parameter (equivalent to the Global Property Support IKE DoS Protection from unidentified Source) to decide what protection to require from remote clients, but / SecureClient clients use the <code>ike_dos_protection</code> . This same client property is called <code>ike_dos_supported_protection_sr</code> on the Security Gateway.	None, Stateless, Puzzles	Puzzles

Protection After Successful Authentication

You can configure fields in Database Tool (GuiDBEdit Tool) (see [sk13009](#)) or dbedit (see [sk13301](#)) to protect against IKE DoS attacks from peers who may authenticate successfully and then attack a Security Gateway. These settings are configured in the Global Properties table and not per Security Gateway. By default these protections are off. Once you enter a value, they will be activated.

To limit the amount of IKE Security Associations (SAs) that a user can open, configure the following fields:

Type of VPN	Field	Recommended Value
Site to site	<code>number_of_ISAKMP_SAs_kept_per_peer</code>	5
Remote user	<code>number_of_ISAKMP_SAs_kept_per_user</code>	5

To limit the amount of tunnels that a user can open per IKE, configure the following fields:

Type of VPN	Field	Recommended Value
Site to site	number_of_ipsec_SAs_per_IKE_SA	30
Remote user	number_of_ipsec_SAs_per_user_IKE_SA	5

Client Properties

Some Security Gateway properties change name when they are downloaded to Remote Access VPN Clients.

The modified name appears in the `userc.C` file, as follows:

Property Name on Security Gateway	Property name on Client in user.C file
ike_dos_protection_unidentified_initiator (Equivalent to the Global Property Support IKE DoS Protection from unidentified Source)	ike_dos_protection or ike_support_dos_protection
ike_dos_supported_protection_sr	ike_dos_protection
ike_dos_puzzle_level_unidentified_initiator	ike_dos_acceptable_puzzle_level
ike_dos_max_puzzle_time_sr	ike_dos_max_puzzle_time

Configuring Advanced IKE Properties

IKE is configured in two places:

- On the VPN community network object (for IKE properties).
- On the Security Gateway network object (for subnet key exchange).

VPN Community Object - Encryption Settings

IPv6 automatically works with IKEv2 encryption only. The option that you select here, applies to IPv4 traffic.

To configure a VPN Community object

1. In SmartConsole, click **Objects** menu > **Object Explorer** (or press Ctrl+E).
2. From the left navigation tree, click **VPN Communities**.

3. Double-click the VPN Community object.

The Community object window opens and shows the **Gateways** page.

4. From the navigation tree, click **Encryption**.
5. Configure the settings.
6. Click **OK**.
7. Install the Access Control Policy.

Encryption Method - for IKE Phase I and IKE Phase II

- **IKEv2 only** - Only support encryption with IKEv2. Security Gateways in this community cannot access peer Security Gateways that support IKEv1 only.
- **Prefer IKEv2, support IKEv1** - If a peer supports IKEv2, the Security Gateway will use IKEv2. If not, it will use IKEv1 encryption. This is recommended if you have a community of older and new Check Point Security Gateways.
- **IKEv1 only** - IKEv2 is not supported.

Encryption Suite

- **Use this encryption suite** - Select the methods negotiated in IKE phase 2 and used in IPSec connections. Select and choose the option for best interoperability with other vendors in your environment.
 - **VPN-A or VPN B** - See RFC 4308 for more information.
 - **Suite-B GCM-128 or 256** - See RFC 6379 for more information.
- **Custom encryption suite** - If you require algorithms other than those specified in the other options, select the properties for IKE Phase 1, including which **Diffie-Hellman group** to use. Also, select properties for IKE Phase 2.

If there is a Security Gateway with Dynamically Assigned IP address inside the VPN community, then R77.30 (or lower) community member Security Gateways that respond to its IKE negotiation, use the configuration defined in SmartConsole > **Menu > Global properties > Remote Access > VPN -Authentication and Encryption**.

More

- **Use aggressive mode** (Main mode is the default) - Select only if the peer only supports aggressive mode. This is only supported with IKEv1.
- **Use Perfect Forward Secrecy**, and the **Diffie-Hellman group** - Select if you need extremely high security.

- **Support IP compression** - Select to decrease bandwidth consumption and for interoperability with third party peers configured to use IP Compression.

VPN Community Object - Advanced Settings

Configure these options in the VPN Community object **Advanced** page:

IKE (Phase 1)

When to renegotiate the IKE Security Associations.

IKE (Phase 2)

When to renegotiate the IPsec security associations. This sets the expiration time of the IPsec encryption keys.

NAT

Disable NAT inside the VPN community - Select to not apply NAT for the traffic while it passes through IPsec tunnels in the community.

Reset

Reset all VPN properties to the default.

Instructions

1. On the **IPsec VPN > VPN Advanced** page, select one of the options in the **VPN Tunnel Sharing** section. There are several settings that control the number of VPN tunnels between peer gateways:

Note - Wire Mode is not supported for IPv6 connections.

- **Use the community settings** - Create the number of VPN tunnels as defined on the community **Tunnel Management** page.

■ **Custom settings:**

- **One VPN tunnel per each pair of hosts** - A VPN tunnel is created for every session initiated between every pair of hosts.
- **One VPN tunnel per subnet pair** - After a VPN tunnel has been opened between two subnets, subsequent sessions between the same subnets will share the same VPN tunnel. This is the default setting and is compliant with the IPsec industry standard.
- **One VPN tunnel per Gateway pair** - One VPN tunnel is created between peer gateways and shared by all hosts behind each peer gateway.

2. On the **Capacity Optimization** page, select limit **Maximum concurrent IKE negotiations**, so you can maximize VPN throughput.

If you have many employees working remotely, you may want to raise the default values.

Link Selection

Link Selection Overview

Link Selection is a method to define which interface is used for incoming and outgoing VPN traffic as well as the best possible path for the traffic. With the Link Selection mechanisms, the administrator can choose which IP addresses are used for VPN traffic on each Security Gateway.

Link Selection has many configuration options to enable you to control VPN traffic. These options include:

- Use probing to choose links according to their availability.
- Use Load Sharing for Link Selection to distribute VPN traffic over available links.
- Use Service Based Link Selection to control bandwidth use.

Configuration settings for remote access clients can be configured together or separately from the Site-to-Site configuration.

Configuring IP Selection by Remote Peer

There are several ways to configure how a Remote Peer resolves the IP address of the local Security Gateway.

You configure the settings in SmartConsole

1. From the left navigation panel, click **Gateways & Servers**.
2. Double-click the Security Gateway object.
3. Click **IPsec VPN > Link Selection**.

Remote peers can connect to the local Security Gateway with one of these settings:

- Always use this IP Address
- Calculate IP based on network topology
- Using DNS resolving
- Using probing - Link redundancy mode

Last Known Available Peer IP Address

The IP address used by a Security Gateway during a successful IKE negotiation with a peer Security Gateway, is used by the peer Security Gateway as the destination IP address for the next IPsec traffic and IKE negotiations that it initiates. This is only the case when the Link Selection configuration does not use probing.

Configuring Outgoing Route Selection

For outbound traffic, there are different methods that can be used to determine which path to use when connecting with a remote peer. These settings are configured in **Security Gateway Properties > IPsec VPN > Link Selection**.

When Initiating a Tunnel:

- **Operating system routing table** (default setting) - With this method, the routing table is consulted for the available route with the lowest metric and best match for the VPN tunnel traffic.
- **Route based probing** - This method also consults the routing table for an available route with the lowest metric and best match. However, before a route is chosen, it is tested for availability with RDP probing. The Security Gateway then selects the best match (highest prefix length) active route with the lowest *metric*. This method is recommended when there is more than one external interface.

If you selected the **IP Selection by Remote Peer** setting of **Use probing** with **Load Sharing**, it also affects **Route based probing** link selection. In this case, **Route based probing** distributes the outgoing encrypted traffic among all available links. All possible links to the peer Security Gateway are derived from the routing table and the link's availability is tested with RDP probing. Every new connection ready for encryption uses the next available link in a round robin manner.

Route based probing enables the use of **On Demand Links (ODL)**, which are triggered upon failure of all primary links. You can run a script to activate an **On Demand Link** when all other links with higher priorities become unavailable. For more information, see the section "**On Demand Links**".

For IKE and RDP sessions, Route based probing uses the same IP address and interface for responding traffic.



Notes:

The High Availability mechanism is based on:

- `resolver_session_interval` (30) - Defines for how many seconds the remote peer status (up or down) stays valid
- `resolver_ttl` (10) - Defines how many seconds the Security Gateway waits before it decides that a remote peer is down

Some network protocols (for example, TCP) might timeout in the time between link failure and the next attempt to resolve. Administrators can decrease these default values. Note that high resolution frequency can overload the Security Gateway. This configuration also changes the default resolution timeouts for the MEP mechanism.

For Layer 2 links, there must be routes to the peer's encryption domains through the local Layer 2 interface device.

When Responding to a Remotely Initiated Tunnel

When responding to a remotely initiated tunnel, there are two options for selecting the interface and next hop that are used. *These settings are only applicable for IKE and RDP sessions.*


These settings are configured in **Link Selection > Outgoing Route Selection > Setup > Link Selection - Responding Traffic** window.

- **Use outgoing traffic configuration** - Select this option to choose an interface with the same method selected in the **Outgoing Route Selection** section of the **Link Selection** page.
- **Reply from the same interface** - This option sends the returning traffic through the same interface and next hop IP address through which it arrived.

Note - When **Route Based Probing** is enabled, **Reply from the same interface** is the selected method and cannot be changed.

Using Route Based Probing

The local Security Gateway, with RDP probing, considers all possible routes between itself and the remote peer Security Gateway. The Security Gateway then decides on the most effective route between the two Security Gateways:



In this scenario, Security Gateway A has two external interfaces, 192.168.10.10 and 192.168.20.10. Peer Security Gateway B also has two external interfaces: 192.168.30.10 and 192.168.40.10.

For Security Gateway A, the routing table reads:

Destination	Netmask	Next hop	Metric
192.168.40.10	255.255.255.0	192.168.10.20	1
192.168.40.10	255.255.255.0	192.168.20.20	2
192.168.30.10	255.255.255.0	192.168.10.20	3
192.168.30.10	255.255.255.0	192.168.20.20	4

For Security Gateway B, the routing table reads:

Destination	Netmask	Next hop	Metric
192.168.20.10	255.255.255.0	192.168.40.20	1
192.168.20.10	255.255.255.0	192.168.30.20	2
192.168.10.10	255.255.255.0	192.168.40.20	3
192.168.10.10	255.255.255.0	192.168.30.20	4

If all routes for outgoing traffic from Security Gateway A are available, the route from 192.168.10.10 to 192.168.40.10 has the lowest metric (highest priority) and is therefore the preferred route.

Source IP Address Settings

Configuration

The source IP address used for outgoing packets can be configured for sessions initiated by the Security Gateway. You configure these settings in **Security Gateway Properties > IPsec VPN > Link Selection > Outgoing Route Selection > Source IP address settings**.

When initiating a VPN tunnel, set the source IP address with one of the following:

- **Automatic (derived from the method of IP selection by remote peer)** - The source IP address of outgoing traffic is derived from the method selected in the **IP Selection by Remote Peer** section.

- If **Main address** or **Selected address from topology table** are chosen in the **IP Selection by Remote Peer** section, then the source IP when initiating a VPN tunnel is the IP specified for that method.
 - If **Calculate IP based on network topology**, **Statically NATed IP**, **Use DNS resolving**, or **Use probing** is chosen in the **IP Selection by Remote Peer** section, then the source IP when initiating a VPN tunnel is the IP address of the chosen outgoing interface.
- **Manual:**
- **Main IP address** - The source IP is derived from the **General Properties** page of the Security Gateway.
 - **Selected address from topology table** - The selected IP from the drop down menu becomes the source IP.
 - **IP address of chosen interface** - The source IP is the same IP of the interface where the traffic is being routed through.

These settings are applicable for RDP and IKE sessions. When responding to an IKE session, use the `reply_from_same_IP` (default: `true`) attribute to follow the settings in the **Source IP address settings** window or to respond from the same IP address.

 **Note** - When Route Based Probing is enabled, `reply_from_same_IP` will be seen as `true`.

Outgoing Link Tracking

When **Outgoing link tracking** is activated on the local Security Gateway, the Security Gateway sends a log for every new resolving decision performed with one of its remote VPN peers.


If **Use Probing** is configured on the local Security Gateway for Remote Peer resolving, or if Route Based Probing is activated on the local Security Gateway, log entries are also created for all resolving changes. For example, if a link in use becomes unavailable and a new available link is chosen, a log entry is issued.

Link Selection Scenarios

Link Selection can be used in many environments. This section describes various scenarios and how Link Selection should be configured in each scenario.

Security Gateway with a Single External Interface

This is the simplest scenario, where the local Security Gateway has a single external interface for VPN:




How do peer Security Gateways select an IP address on the local Security Gateway for VPN traffic?

Since there is only one interface available for VPN, to determine how remote peers determine the IP address of the local Security Gateway, select the following from the **IP Selection by Remote Peer** section of the Link Selection page:

- Select **Main address** or choose an IP address from the **Selected address from topology table** drop down menu.
- If the IP address is located behind a static NAT device, select **Statically NATed IP**.

Security Gateway with Several IP Addresses Used by Different Parties

In this scenario, the local Security Gateway has a point-to-point connection from two different interfaces. Each interface is used by a different remote party:




The local Security Gateway has two IP addresses used for VPN. One interface is used for VPN with a peer Security Gateway A and one interface for peer Security Gateway B.

To determine how peer Security Gateways discover the IP address of the local Security Gateway, enable **one-time probing** with **High Availability** redundancy mode. Since only one IP is available for each peer Security Gateway, probing only has to take place one time.

Security Gateway with an Interface Behind a Static NAT Device

In this scenario, the local Security Gateway has two external interfaces available for VPN.

The IP address of interface `eth0` is translated using a NAT device:



To determine how peer Security Gateways discover the IP address of the local Security Gateway, use **ongoing probing** with **High Availability** redundancy mode.


In order for the Static NAT IP address to be probed, it must be added to the **Probe the following addresses** list in the **Probing Settings** window.

Utilizing Load Sharing

Depending on your configuration, there are many ways to use Load Sharing to distribute VPN traffic among available links between the local and peer Security Gateways.

Load Sharing with Multiple External Interfaces on Each End

In the following scenario, the local and peer Security Gateways each have two external interfaces available for VPN traffic.




To utilize both external interfaces by distributing VPN traffic among all available links, use the Probing redundancy mode of **Load Sharing** on both Security Gateways. You can also specify that only certain external interfaces should be probed by putting only those interfaces in the **Probe the following addresses** list in the **Probing Settings** window. If one link goes down, traffic will automatically be rerouted through the other link.

To enable this configuration, make sure that your routing table allows packet flow back and forth between both `eth0` interfaces and packet flow back and forth between both `eth1` interfaces. Then Link Selection can reroute the VPN traffic between these available links.

Load Sharing with Multiple External Interfaces on One End

In the following scenario, the local Security Gateway has two external interfaces available for VPN traffic. The peer Security Gateway has one external interface for VPN traffic.



To utilize both external interfaces and distribute VPN traffic between the available links, use the **Probing redundancy mode of Load Sharing** on the local Security Gateway. Then the peer Security Gateway will distribute its outgoing VPN traffic between interfaces `eth0` and `eth1` of the local Security Gateway.

If the default, **Operating system routing table**, setting in the **Outgoing Route Selection** section is selected, the local Security Gateway will only use one of its local interfaces for outgoing VPN traffic; the route with the lowest metric and best match to reach the single IP address of the peer Security Gateway, according to the routing table.

If you want to distribute the outgoing VPN traffic on both outbound links from the local Security Gateway as well, select **Route Based Probing** in the **Outgoing Route Selection** on the **Link Selection** page of the local Security Gateway.

Service Based Link Selection

Service Based Link Selection enables administrators to control outgoing VPN traffic and bandwidth use by assigning a service or a group of services to a specific interface for outgoing VPN routing decisions. The encrypted traffic of an outgoing connection is routed through the configured interface according to the traffic's service. The links to the peer Security Gateway are derived from the routing table and the link's availability is tested with RDP probing.

If all links through the interface assigned to a specific service stop responding to RDP probing, a link failover will occur by default, as in any other probing mode. When a link through the assigned interface is restored, new outgoing connections are assigned to it, while existing connections are maintained over the backup link until they are completed.

It is possible to configure the traffic of a specific service not to fail over. In this case, traffic of the configured service will only be routed through interfaces assigned to this service, even if these interfaces stop responding to RDP.

If the same service is assigned to more than one interface, this service's traffic is distributed between the configured interfaces. Every new outgoing encrypted connection uses the next available link in a round robin manner.

All traffic from services that are not assigned to a specific interface is distributed among the remaining interfaces. If all links through these interfaces are down, the traffic is distributed among the interfaces that are configured for specific services.

Service Based Link Selection configuration requires enabling the following features:

- **IP Selection by Remote Peer** - Load Sharing probing mode
- **Outgoing Route Selection** - Route based probing
- **Service Based Link Selection** - Configuration file on the management server

Service Based Link Selection is supported on Security Gateways R71 and higher.

Configuring Service Based Link Selection

To configure Service Based Link Selection:

1. In the **Link Selection** page, in the **IP Selection by Remote Peer** section, select:
 - **Use probing. Redundancy mode**
 - **Load Sharing**
2. In the **Outgoing Route Selection** section, select **Route based probing**.

Edit the Service Based Link Selection configuration in the `$FWDIR/conf/vpn_service_based_routing.conf` configuration file on the management server.

Fill in each line in the configuration file to specify the target Security Gateway, the interface for outgoing routing, and the service (or services group) to route through this interface. Use the names defined in the SmartConsole. Fill in all of the details for each Security Gateway on which you want to configure Service Based Link Selection.

The fields are:


- **Gateway** - Security Gateway name (the name of the VPN Security Gateway or the cluster).
- **Interface** - Interface name.
- **Service** - Service or services group name.
- **dont_failover** - (Optional) If this string is present, traffic of the configured service will only be routed through interfaces configured for this service and will not fail over to another interface.

Service Based Link Selection Scenarios

The following scenarios provide examples of how Service Based Link Selection can be utilized.

Service Based Link Selection with Two Interfaces on Each End

In the scenario below, the local and peer Security Gateways each have two external interfaces for VPN traffic.



In this example, interface `eth1` of both Security Gateways is dedicated to HTTP and FTP traffic. All other traffic is routed to interface `eth0`.

If the available link through `eth1` stops responding to RDP probing, HTTP and FTP traffic will fail over to `eth0`. It is possible to specify that HTTP and FTP traffic should only be routed through `eth1` even if the link through `eth1` stops responding. Specify this by including the **`dont_failover`** flag when editing the Service Based Link Selection configuration file.

All other traffic that is not HTTP or FTP will be routed through `eth0`. If the link through `eth0` stops responding to RDP probing, all traffic will be routed through `eth1`.

The Service Based Link Selection configuration file for this environment should appear as follows:


Security Gateway	Interface	Service	[dont_failover]
London_GW	eth1	http	
London_GW	eth1	ftp	
Paris_GW	eth1	http	
Paris_GW	eth1	ftp	

Alternatively, in SmartConsole, you can create a Services Group that includes HTTP and FTP services. In the example below, this group is called **http_ftp_grp**. With this group, the Service Based Link Selection configuration file for this environment should appear as follows:

Security Gateway	Interface	Service	[dont_failover]
London_GW	eth1	http_ftp_grp	
Paris_GW	eth1	http_ftp_grp	

Service Based Link Selection with Multiple Interfaces on Each End

In the following scenario, the local and peer Security Gateways each have three external interfaces available for VPN.




To utilize all three external interfaces and distribute the VPN traffic among the available links, Link Selection Load Sharing and Route based probing should be enabled. To control your bandwidth use, dedicate one or more links to a specific service or services using Service Based Link Selection. In this scenario, interfaces `eth0` and `eth1` of both Security Gateways are dedicated to SIP traffic. SIP traffic is distributed between `eth0` and `eth1`. All other traffic is routed through `eth2`.

If either the link through `eth0` or the link through `eth1` stops responding to RDP probing, SIP traffic will fail over to the other SIP interface. If the link through `eth2` stops responding to RDP probing, all traffic will be routed through `eth0` or `eth1`.

Service Based Link Selection with Two Interfaces on One End

In the following scenario, the local Security Gateway has two external interfaces available for VPN traffic. The peer Security Gateway has a single external interface for VPN traffic.



To utilize all external interfaces and distribute the VPN traffic among the available links, Link Selection Load Sharing and Route based probing should be enabled on the local Security Gateway, **London_GW**. To control your bandwidth use, dedicate interface `eth1` of the local Security Gateway to HTTP and FTP traffic using Service Based Link Selection. The local Security Gateway will route outgoing HTTP and FTP connections through interface `eth1`. All other traffic, not HTTP or FTP, will be routed through `eth0`.

In this scenario, HTTP and FTP traffic should not fail over. HTTP and FTP traffic should only be routed through interface `eth1`, even if the link through interface `eth1` stops responding to RDP probing. This is configured by specifying the **dont_failover** flag.

The Service Based Link Selection configuration file for this environment should appear as follows:

Security Gateway	Interface	Service	[dont_failover]
London_GW	eth1	http	dont_failover
London_GW	eth1	ftp	dont_failover

Since the Service Based Link Selection configuration is only applicable for outgoing traffic of the local Security Gateway, the peer Security Gateway can send HTTP and FTP traffic to either interface of the local Security Gateway. The outgoing VPN traffic of the peer Security Gateway is distributed between interfaces `eth0` and `eth1` of the local Security Gateway.

Trusted Links

Trusted Links allows you to set an interface as "trusted" for VPN traffic so that traffic sent on that link will not be encrypted. You may want to set up a trusted link if you are confident that the link is already encrypted and secure and you do not need a second encryption.

If you configure an interface as trusted, traffic routed through that interface will be sent unencrypted, while traffic sent through other interfaces will still be encrypted.

Trusted interfaces should be configured symmetrically on the local and peer Security Gateways. If only one side of the link is configured as trusted for VPN traffic, clear traffic received by a non-trusted interface will be dropped by the peer Security Gateway.

If you have configured a specific link as trusted for VPN traffic and you use probing, the probing method considers all links, including the trusted link, when choosing a link for a connection.

The probing method chooses the link according to these criteria:

- The configured redundancy mode, High Availability or Load Sharing
- If Service Based Link Selection is configured.

If the trusted link is chosen for a connection, the traffic is not encrypted. If another, non-trusted, link is chosen, the traffic is encrypted.

In an MEP configuration, trusted links are only supported for connections initiated by a peer Security Gateway to a MEP Security Gateway. Unencrypted VPN connections routed through a trusted interface and initiated by a MEP Security Gateway may be dropped by the peer Security Gateway.

Trusted links are not supported in Traditional mode. In Traditional mode, trusted link settings are ignored and VPN traffic is always encrypted.

Trusted links are supported on Security Gateways R71 and higher.

Configuring Trusted Links


Use the Database Tool (GuiDBEdit Tool) (see [sk13009](#)) to configure Trusted Links.

To configure a trusted link:

1. In the top left pane, go to **Network objects > network_objects**.
2. In the top right pane, click the Security Gateway / Cluster object that you want to edit.
3. In the bottom pane, search for the interface that you want to configure as trusted from within the interfaces set. The interface name appears in the `officialname` attribute.
4. Within the trusted interface set, change the value of the `vpn_trusted` attribute to `true` (default value: `false`).
5. Configure trusted interfaces symmetrically on the peer Security Gateways. If only one side of the link is configured as trusted for VPN traffic, clear traffic received by a non-trusted interface will be dropped by the peer Security Gateway.
6. Save changes (**File** menu > **Save All**).
7. In SmartConsole, install the Access Control Policy on the Security Gateway / Cluster object.

Trusted Links Scenarios

In the following scenario, both the local and peer Security Gateways have two external interfaces available for VPN traffic. Interface `eth1` on both Security Gateways has been configured as a trusted interface. Therefore, traffic sent from `eth1` of the local Security Gateway is sent unencrypted and is accepted by interface `eth1` of the peer Security Gateway, and the other way around.




If the probing redundancy mode is High Availability and the trusted link is configured as the **Primary IP address**, the trusted link will be used for VPN traffic. If the trusted link stops responding to RDP probing, the link through Interface `eth0` will be used for VPN traffic and traffic will be encrypted.

If the probing redundancy mode is Load Sharing, the VPN traffic will be distributed between the available links. Connections routed through interface `eth0` will be encrypted while connections routed through the trusted link will not be encrypted.

Using Trusted Links with Service Based Link Selection

In the following scenario, the local and peer Security Gateways have two external interfaces available for VPN traffic. Interface `eth1` on both Security Gateways is configured as a trusted interface for VPN traffic since encryption is not needed on that link. In addition, interface `eth1` of both Security Gateways is dedicated to SIP traffic using Service Based Link Selection.




SIP traffic is routed through the trusted link between the two `eth1` interfaces and will not be encrypted. If the trusted link stops responding to RDP probing, SIP traffic will be routed through the `eth0` interfaces and will be encrypted.

All other traffic that is not SIP is encrypted and routed through the interface `eth0` link. However, if interface `eth0` stops responding to RDP probing, all the traffic will be routed through the trusted link and will not be encrypted.

On Demand Links (ODL)

Route based probing enables use of an On Demand Link (ODL), which is triggered upon failure of all primary links. When a failure is detected, a custom script is used to activate the ODL and change the applicable routing information. The ODL's metric must be set to be larger than a configured minimum in order for it to be considered an ODL.



The Security Gateway has two external links for Internet connectivity: one to an ISP, the other to an ISDN dialup. The ISDN dialup connection is configured as an On Demand Link.

On the Security Gateway, the Route Based Probing mechanism probes all of the non-On Demand Links and selects the active link with the lowest metric. In this case, it probed the ISP link. A script is run to activate the On Demand Link when all other links with higher priorities become unavailable. When the link becomes available again, a shutdown script is run automatically and the connection continues through the link with the ISP.

Note - On Demand Links are probed only once with a single RDP session. Fail over between On Demand Links is not supported.

Configuring On Demand Links

You can enable On Demand Links only if you enabled Route Based Probing. Configure On Demand Links commands in Database Tool (GuiDBEdit Tool) (see [sk13009](#)).

Property	Description
<code>use_on_demand_links</code>	Enables on-demand links. The default is <code>false</code> . Change to <code>true</code> .
<code>on_demand_metric_min</code>	Defines the minimum metric level for an on-demand link. This value must be equal to or higher than the configured minimum metric.

Property	Description
<code>on_demand_initial_script</code>	The name of the on-demand script, which runs when all not-on-demand routes stop responding. Put the script in the <code>\$FWDIR/conf/</code> directory.
<code>on_demand_shutdown_script</code>	This script is run when the failed links become available. Put the script in the <code>\$FWDIR/conf/</code> directory.

If you do not want to use Database Tool (GuiDBEdit Tool), you can configure the `use_on_demand_links` and `on_demand_metric_min` settings in SmartConsole:

1. Click **Menu > Global properties**.
2. Click **Advanced > Configure**.
3. Click **VPN Advanced Properties > Link Selection**.
4. Select `use_on_demand_links` to enable On Demand Links.
5. In the `on_demand_metric_min` field, set the minimum metric level for an On Demand Link.
6. Click **OK**.
7. Click **OK**.
8. Install the Access Control Policy.

Link Selection and ISP Redundancy

ISP Redundancy enables reliable Internet connectivity by allowing a single or clustered Security Gateway to connect to the Internet via redundant ISP connections. As part of standard VPN installation, it offers two modes of operation:

- Load Sharing mode
- Primary/Backup mode

Configuring Link Selection and ISP Redundancy

Configure Link Selection and ISP Redundancy in the **Other > ISP Redundancy** page of the Security Gateway object:

- **Load Sharing** mode connects to both ISPs while sharing the load of outgoing connections between the ISPs according to a designated weight assignment. New connections are randomly assigned to a link. If a link fails, all new outgoing connections are directed to the active link. This configuration effectively increases the WAN bandwidth while providing connectivity protection. The assigned ISP Links weight is only supported for Security Gateway traffic.
- **Primary/Backup** mode connects to an ISP through the primary link, and switches to a backup ISP if the primary ISP link fails. When the primary link is restored, new outgoing connections are assigned to it, while existing connections are maintained over the backup link until they are complete.


The settings configured in the **ISP Redundancy** window are by default, applied to the **Link Selection** page and will overwrite any pre-existing configuration. The following settings carry over:

- When ISP Redundancy is configured, the default setting in the Link Selection page is **Use ongoing probing**. However, Link Selection only probes the ISPs configured in the **ISP Redundancy** window. This enables connection failover of the VPN tunnel if connectivity to one of the Security Gateway interfaces fails.
- If the ISP Redundancy mode is **Load Sharing**, the Probing redundancy mode in the Link Selection page is also **Load Sharing**.
- If the ISP Redundancy mode is **Primary/Backup**, the Probing redundancy mode in the Link Selection page is **High Availability**.
 - The Primary ISP link of the ISP redundancy is set as the Primary Address of the Link Selection probing. The Primary Address is set under: **IP Selection by Remote Peer > Use Probing > Configure** (or **View**, if the settings are derived from the ISP Redundancy settings).

If you do not want the ISP Redundancy settings to affect the Link Selection settings, on the ISP Redundancy page, clear the check box that says **Apply settings to VPN traffic** and configure the required VPN settings on the **Link Selection** page. This may apply when you want to route VPN traffic differently than the clear traffic. For example, if you want to use Load Sharing for clear traffic and High Availability for VPN traffic, or if you want to use different primary ISPs for clear and VPN traffic.


Link Selection and ISP Redundancy

In the following scenario, the local Security Gateway maintains links to ISPs A and B, both of which provide connectivity to the Internet with ISP Redundancy.



In the **Topology > ISP Redundancy** window, configure the ISP Redundancy settings, such as ISP Links and Redundancy mode. The ISP Redundancy settings are applied by default to VPN traffic. The derived Link Selection settings are visible in the **IPsec VPN > Link Selection** window.

In the following scenario, the **Apply settings to VPN traffic** on the **ISP Redundancy** page was cleared, and there are different settings configured for **Link Selection** and **ISP Redundancy**.



In this scenario:

- Security Gateways A, B, and C each have two interfaces configured as ISP links.
- ISP Redundancy is configured on Security Gateway A.
- Security Gateway A should use ISP 1 in order to connect to Security Gateway B and ISP 2 in order to connect to Security Gateway C. If one of the ISP links becomes unavailable, the other ISP should be used.

In this scenario, the administrator of Security Gateway A needs to:

- Clear the box **Apply settings to VPN traffic** in the **ISP Redundancy** window.
- Reconfigure the **Outgoing Route Selection** to **Route Based Probing** in the **Link Selection** window.
- Configure the routing table so that ISP 1 is the highest priority for peer Security Gateway B and ISP 2 has the highest priority for peer Security Gateway C.

Link Selection with non-Check Point Devices

RDP probing, the probing method used for certain Link Selection features, is proprietary to Check Point and only works between Check Point entities. It is not supported with non-Check Point devices.

Since RDP probing is not active on non-Check Point Security Gateways, the following results apply if a Check Point Security Gateway sends VPN traffic to a non-Check Point Gateway:

- **Use probing** cannot be used by locally managed Check Point Security Gateways to determine the IP address of non-Check Point devices. Any of the other methods available from the **IP Selection by Remote Peer** section can be used.
- **Load Sharing** and **Service Based Link Selection** do not work with non-Check Point Gateways. If Load Sharing or Service Based Link Selection is enabled on the local Security Gateway, but the peer is a non-Check Point device, the local Security Gateway will only use one link to the non-Check Point device: the best match (highest prefix length) link with the lowest metric.
- If **Route based probing** is selected as the **Outgoing Route Selection** method, for VPN traffic to a non-Check Point device, the local Security Gateways always use the best match (highest prefix length) link with the lowest metric.

Public Key Infrastructure

Need for Integration with Different PKI Solutions

X.509-based PKI solutions provide the infrastructure that enables entities to establish trust relationships between each other based on their mutual trust of the Certificate Authority (CA). The trusted CA issues a certificate for an entity, which includes the entity's public key. Peer entities that trust the CA can trust the certificate - because they can verify the CA's signature - and rely on the information in the certificate, the most important of which is the association of the entity with the public key.

IKE standards recommend the use of PKI in VPN environments, where strong authentication is required.

A Security Gateway taking part in VPN tunnel establishment must have an RSA key pair and a certificate issued by a trusted CA. The certificate contains details about the module's identity, its public key, CRL retrieval details, and is signed by the CA.

When two entities try to establish a VPN tunnel, each side supplies its peer with random information signed by its private key and with the certificate that contains the public key. The certificate enables the establishment of a trust relationship between the Security Gateways; each Security Gateway uses the peer Security Gateway public key to verify the source of the signed information and the CA's public key to validate the certificate's authenticity. In other words, the validated certificate is used to authenticate the peer.

Every deployment of Check Point Security Management Server includes an Internal Certificate Authority (ICA) that issues VPN certificates for the VPN modules it manages. These VPN certificates simplify the definition of VPNs between these modules.

Situations can arise when integration with other PKI solutions is required, for example:

- A VPN must be established with a Security Gateway managed by an external Security Management Server. For example, the peer Security Gateway belongs to another organization which utilizes Check Point products, and its certificate is signed by its own Security Management Server ICA.
- A VPN must be established with a non-Check Point VPN entity. In this case, the peer's certificate is signed by a third-party CA.
- An organization may decide, for whatever reason, to use a third party CA to generate certificates for its Security Gateways.

Supporting a Wide Variety of PKI Solutions

Check Point Security Gateways support many different scenarios for integrating PKI in VPN environments:

- **Multiple CA Support for Single VPN Tunnel** - Two Security Gateways present a certificate signed by different ICAs.
- **Support for non-ICA CAs** - In addition to ICA, Security Gateways support the following Certificate Authorities:
 - **External ICA** - The ICA of another Security Management Server
 - **Other OPSEC certified PKI solutions**
- **CA Hierarchy** - CAs are typically arranged in a hierarchical structure where multiple CAs are subordinate to a root authority CA. A subordinate CA is a Certificate Authority certified by another Certificate Authority. Subordinate CAs can issue certificates to other, more subordinate CAs, forming a certification chain or hierarchy.

PKI and Remote Access Users

The Check Point Suite supports certificates not only for Security Gateways but for users as well. For more information, see Introduction to Remote Access VPN for information about user certificates.

PKI Deployments and VPN

Following are some sample CA deployments:


- Simple Deployment - internal CA
- CA of an external Security Management Server
- CA services provided over the Internet
- CA on the LAN

Simple Deployment Internal CA

When the VPN tunnel is established between Security Gateways managed by the same Security Management Server, each peer has a certificate issued by the Security Management Server's ICA.

CA of An External Security Management Server


If a Check Point Security Gateway is managed by an external Security Management Server (for example, when establishing a VPN tunnel with another organization's VPN modules), each peer has a certificate signed by its own Security Management Server's ICA.



Security Management Server A issues certificates for Security Management Server B that issues certificates for Security Gateway B.

CA Services Over the Internet


If the certificate of a Security Gateway is issued by a third party CA accessible over the Internet, CA operations such as registration or revocation are usually performed through HTTP forms. CRLs are retrieved from an HTTP server functioning as a CRL repository.



Security Gateways A and B receive their certificates from a PKI service provider accessible via the web. Certificates issued by external CAs may be used by Security Gateways managed by the same Security Management Server to verification.

CA Located on the LAN

If the peer VPN Security Gateway certificate is issued by a third party CA on the LAN, the CRL is usually retrieved from an internal LDAP server, as shown:



Trusting An External CA

A trust relationship is a crucial prerequisite for establishing a VPN tunnel. However, a trust relationship is possible only if the CA that signs the peer's certificate is "trusted." Trusting a CA means obtaining and validating the CA's own certificate. Once the CA's Certificate has been validated, the details on the CA's certificate and its public key can be used to both obtain and validate other certificates issued by the CA.

The Internal CA (ICA) is automatically trusted by all modules managed by the Security Management Server that employs it. External CAs (even the ICA of another Check Point Security Management Server) are not automatically trusted, so a module must first obtain and validate an external CA's certificate. The external CA must provide a way for its certificate to be imported into the Security Management Server.

If the external CA is:

- The ICA of an external Security Management Server, see the [R81.10 Security Management Administration Guide](#) for further information
- An OPSEC Certified CA, use the CA options on the **Servers and OPSEC Applications** tab to define the CA and obtain its certificate

Subordinate Certificate Authorities

A subordinate CA is a Certificate Authority certified by another Certificate Authority. Subordinate CAs can issue certificates to other, more subordinate CAs, in this way forming a certification chain or hierarchy. The CA at the top of the hierarchy is the root authority or root CA. Child Certificate Authorities of the root CA are referred to as Subordinate Certificate Authorities.

With the CA options on the **Servers and OPSEC Applications** tab, you can define either a Certificate Authority as either Trusted or Subordinate. Subordinate CAs are of the type OPSEC, and not trusted.

Enrolling a Managed Entity

Enrollment means requesting and obtaining a certificate from a CA, for an entity.

The process of enrollment begins with the generation of a key pair. A certificate request is then created out of the public key and additional information about the module. The type of the certificate request and the rest of the enrollment process depends on the CA type.

The case of an internally managed Security Gateway is the simplest, because the ICA is located on the Security Management Server machine. The enrollment process is completed automatically.

To obtain a certificate from an OPSEC Certified CA, Security Management Server takes the module details and the public key and encodes a PKCS#10 request. The request (which can include *SubjectAltName* for OPSEC certificates and Extended Key Usage extensions) is delivered to the CA manually by the administrator. Once the CA issues the certificate the administrator can complete the process by importing the certificate to the Security Management Server.

A certificate can also be obtained for the Security Gateway with Automatic Enrollment. With Automatic Enrollment, you can automatically issue a request for a certificate from a trusted CA for any Security Gateway in the community. Automatic Enrollment supports the following protocols:

- SCEP
- CMPV1
- CMPV2

Note - During SCEP enrollment, some HTTP requests may be larger than 2000 bytes, and may be dropped by the HTTP protocol inspection mechanism if enabled. A change of the default value will be required to enable these HTTP requests. If enrollment still fails, enrollment must be done manually. For more information, see the [R81.10 Threat Prevention Administration Guide](#).

Validation of a Certificate

When an entity receives a certificate from another entity, it must:

1. Verify the certificate signature, i.e. verify that the certificate was signed by a trusted CA. If the certificate is not signed directly by a trusted CA, but rather by a subsidiary of a trusted CA, the path of CA certificates is verified up to the trusted CA.
2. Verify that the certificate chain has not expired.
3. Verify that the certificate chain is not revoked. A CRL is retrieved to confirm that the serial number of the validated certificate is not included among the revoked certificates.

In addition, VPN verifies the validity of the certificate's use in the given situation, confirming that:

- The certificate is authorized to perform the required action. For example, if the private key is needed to sign data (e.g., for authentication) the **KeyUsage** extension on the certificate - if present - is checked to see if this action is permitted.
- The peer used the correct certificate in the negotiation. When creating a VPN tunnel with an externally managed module, the administrator may decide that only a certificate signed by a specific CA from among the trusted CAs can be accepted. (Acceptance of certificates with specific details such as a *Distinguished Name* is possible as well).

Revocation Checking

There are two available methods useful in determining the status of a certificate:

1. CRL
2. Online Certificate Status Protocol (OCSP)

Enrolling with a Certificate Authority

A certificate is automatically issued by the Internal Certificate Authority for all internally managed entities that are VPN-capable. That is, after the administrator enables the **IPsec VPN Software Blade** in a Security Gateway or Cluster object (on the **General Properties** page > on the **Network Security** tab).

The process for obtaining a certificate from an **OPSEC PKI CA** or **External Check Point CA** is identical.

Manual Enrollment with OPSEC Certified PKI

To create a PKCS#10 Certificate Request:

1. Create a Certificate Authority object.
2. From the left navigation panel, click **Gateways & Servers**.
3. Double-click the applicable Security Gateway or Cluster object.
4. From the left tree, click **General Properties** and make sure to enable the **IPsec VPN Software Blade**.
5. From the left tree, click **IPsec VPN**.

6. In the section **Repository of Certificates Available to the Gateway**, click **Add**.

The **Certificate Properties** window opens.

7. In the **Certificate Nickname** field, enter a text string.

The nickname is only an identifier and has no bearing on the content of the certificate.

8. From the drop-down menu **CA to enroll from**, select the Certificate Authority that issues the certificate.

Note - The menu shows only trusted Certificate Authorities and subordinate Certificate Authorities that lead directly to a trusted Certificate Authority. If the CA that issues the certificate is a subordinate CA that does not lead directly to a trusted CA, it is not in the menu.

9. In the section **Key pair generation and storage**, select the applicable method:

- **Store keys on the Security Management server** - Certificate creation is performed entirely between the Management Server and applicable CA. The keys and the certificate are downloaded securely to the Security Gateway (Cluster Members) during policy installation.
- **Store keys on the Module** - Management Server directs the Security Gateway (or Cluster Members) to create the keys and supply only the required material for creation of the certificate request. Only the certificate is downloaded to the Security Gateway (Cluster Members) during policy installation.

10. Click **Generate**.

The **Generate Certificate Properties** window opens.

11. Enter the applicable DN.

The CA administrator determines the final DN that appears in the certificate.

If a **Subject Alternate Name** extension is required in the certificate, select **Define Alternate Name**.

The public key and the DN are then used to DER-encode a PKCS#10 Certificate Request.

-  **Note** - Adding the object's IP address as the **Alternate Name** extension can be configured as a default setting.

This configuration also applies for Internal Certificate Authorities.

- a. In SmartConsole, click **Menu > Global properties > Advanced > Configure**.
 - b. Click **Certificates and PKI properties**.
 - c. Select these options:
 - **add_ip_alt_name_for_ICA_certs** (closer to the top of this page)
 - **add_ip_alt_name_for_opsec_certs** (closer to the bottom of this page)
 - d. Click **OK** to close the **Advanced Configuration** window.
 - e. Click **OK** to close the **Global properties** window.
12. When the certificate appears in the section **Repository of Certificates Available to the Gateway**:
- a. Select this certificate.
 - b. Click **View**.
 - c. In the **Certificate View** window:
 - i. Click inside the window.
 - ii. Select the whole text (press the **CTRL+A** keys, or right-click the mouse and click **Select All**).
 - iii. Copy the whole text (press the **CTRL+C** keys, or right-click the mouse and click **Copy**).
 - iv. Paste the text into a plain text editor (like Notepad).
 - v. Click **OK**.
13. Send the certificate information to the Certificate Authority administrator.
- The CA administrator must now complete the task of issuing the certificate.
- Different CAs provide different ways of doing this, such as an advanced enrollment form (as opposed to the regular form for users).
- The issued certificate may be delivered in various ways, for example, email.
14. After the certificate arrives from the Certificate Authority administrator, you must save it in the Certificate Authority object:
- a. In SmartConsole, click **Objects > Object Explorer** (or press the **CTRL+E** keys).
 - b. In the left tree, click **Servers**.

- c. Double-click the applicable Certificate Authority object.
- d. Click the **OPEC PKI** tab.
- e. In the **Certificate** section, click **Get**.
- f. Browse to the location, where you saved the certificate file.
- g. Select the certificate file and click **Open**.
- h. If the certificate details are correct, click **OK** to accept this certificate.
- i. Click **OK** to close the **Certificate Authority Properties** window.
- j. Close the **Object Explorer** window.

15. Publish the SmartConsole session

Automatic Enrollment with the Certificate Authority

On the **OPSEC PKI** tab of the Certificate Authority object:

1. Select the option **Automatically enroll certificate**.
2. Select the applicable protocol - **scep** or **cmp**.

Follow these steps:

1. From the left navigation panel, click **Gateways & Servers**.
2. Double-click the applicable Security Gateway or Cluster object.
3. From the left tree, click **General Properties** and make sure to enable the **IPsec VPN Software Blade**.
4. From the left tree, click **IPsec VPN**.
5. In the section **Repository of Certificates Available to the Gateway**, click **Add**.

The **Certificate Properties** window opens.

6. In the **Certificate Nickname** field, enter a text string.

The nickname is only an identifier and has no bearing on the content of the certificate.

7. From the drop-down menu **CA to enroll from**, select the Certificate Authority that issues the certificate.

Note - The menu shows only trusted CAs and subordinate CAs that lead directly to a trusted CA. If the CA that issues the certificate is a subordinate CA that does not lead directly to a trusted CA, it is not in the menu.

8. In the section **Key pair generation and storage**, select the applicable method:

- **Store keys on the Security Management server** - Certificate creation is performed entirely between the Management Server and applicable CA. The keys and the certificate are downloaded securely to the Security Gateway (Cluster Members) during policy installation.
 - **Store keys on the Module** - Management Server directs the Security Gateway (or Cluster Members) to create the keys and supply only the required material for creation of the certificate request. Only the certificate is downloaded to the Security Gateway (Cluster Members) during policy installation.
9. Click **Generate** and select **Automatic enrollment**.
The **Generate Keys and Get Automatic Enrollment Certificate** window opens.
 - Supply the **Key Identifier** and your secret **Authorization code**.
 - Click **OK**.
 10. When the certificate appears in the section **Repository of Certificates Available to the Gateway**:
 - a. Select this certificate.
 - b. Click **View**.
 - c. In the **Certificate View** window, click **Copy to Clipboard** or **Save to File**.
 11. Send the request to CA administrator.
Different Certificate Authorities provide different means for doing this. For example, an advanced enrollment form on their website. The issued certificate can be delivered in various ways, such as by email. After you receive the certificate, save it to disk.
 12. From the left tree click, **IPsec VPN**.
 13. In the section **Repository of Certificates Available to the Gateway**:
 - a. Select the applicable certificate.
 - b. Click **Complete**.
 14. Browse to the folder where you stored the issued certificate, select the certificate, and examine the certificate details.
 15. Click **OK** to close the Security Gateway or Cluster object.
 16. Publish the SmartConsole session

Enrolling through a Subordinate CA

When enrolling through a Subordinate CA:

- Supply the password of the Subordinate CA which issues the certificate (not the CA at the top of the hierarchy).
- The Subordinate CA must lead directly to a trusted CA.

CRL

VPN can retrieve the CRL from either an HTTP server or an LDAP server. If the CRL repository is an HTTP server, the module uses the URL published in the CRL **Distribution Point** extension on the certificate and opens an HTTP connection to the CRL repository to retrieve the CRL.

If the CRL repository is an LDAP server, VPN attempts to locate the CRL in one of the defined LDAP account units. In this scenario, an LDAP account unit must be defined. If the CRL **Distribution Point** extension exists, it publishes the DN of the CRL, namely, the entry in the Directory under which the CRL is published or the LDAP URI. If the extension does not exist, VPN attempts to locate the CRL in the entry of the CA itself in the LDAP server.

OCSP

Online Certificate Status Protocol (OCSP) enables applications to identify the state of a certificate. OCSP may be used for more timely revocation information than is possible with CRLs and may also be used to obtain additional status information. When OCSP client issues a status request to an OCSP server, acceptance of the certificate in question is suspended until the server provides a response.

In order to use OCSP, the root CA must be configured to use this method instead of CRL. This setting is inherited by the subordinate CAs.

CRL Prefetch-Cache

Since the retrieval of CRL can take a long time (in comparison to the entire IKE negotiation process), VPN stores the CRLs in a CRL cache so that later IKE negotiations do not require repeated CRL retrievals.

The cache is pre-fetched:

- every two hours
- on policy installation
- when the cache expires

If the pre-fetch fails, the previous cache is not erased.

Note - The ICA requires the use of a CRL cache.

An administrator can shorten the lifetime of a CRL in the cache or even to cancel the use of the cache. If the CRL Cache operation is canceled, the CRL must be retrieved for each subsequent IKE negotiation, thus considerably slowing the establishment of the VPN tunnel. Because of these performance implications, we recommend that you only disable CRL caching when the level of security demands continuous CRL retrieval.

Special Considerations for the CRL Pre-fetch Mechanism

The CRL pre-fetch mechanism makes a "best effort" to obtain the most up to date list of revoked certificates. However, after the `cpstop` and `cpstart` commands have been executed, the cache is no longer updated. The Security Gateway continues to use the old CRL for as long as the old CRL remains valid (even if there is an updated CRL available on the CA). The pre-fetch cache mechanism returns to normal functioning only after the old CRL expires and a new CRL is retrieved from the CA.

In case there is a requirement that after the `cpstop` and `cpstart` commands, the CRLs will be updated immediately, proceed as follows:

- After executing the "`cprestart`" command, run the ["vpn_crl_zap" on page 214](#) command to empty the cache, or:
- In SmartConsole
 1. Click **Menu > Global properties > Advanced > Configure**.
 2. Click **Certificates and PKI properties**.
 3. Select **flush_crl_cache_file_on_install**.
 4. Click **OK**.
 5. Install the Access Control Policy.

When a new policy is installed, the cache is flushed, and a new CRL will be retrieved on demand.

CRL Grace Period

Temporary loss of connection with the CRL repository or slight differences between clocks on the different machines may cause valid of CRLs to be considered invalid - and thus the certificates to be invalid as well. VPN overcomes this problem by supplying a CRL Grace Period. During this period, a CRL is considered valid even if it is not valid according to the CRL validity time.

Special Considerations for PKI

Using the Internal CA vs. Deploying a Third Party CA

The Internal CA makes it easy to use PKI for Check Point applications such as site-to-site and remote access VPNs. However, an administrator may prefer to continue using a CA that is already functioning within the organization, for example a CA used to provide secure email, and disk encryption.

Distributed Key Management and Storage

Distributed Key Management (DKM) provides an additional layer of security during the key generation phase. Instead of the Security Management Server generating both public and private keys and downloading them to the module during a policy installation, the management server instructs the module to create its own public and private keys and send (to the management server) only its public key. The private key is created and stored on the module in either a hardware storage device, or via software that emulates hardware storage. Security Management Server then performs certificate enrollment. During a policy installation, the certificate is downloaded to the module. The private key never leaves the module.

Local key storage is supported for all CA types.

DKM is supported for all enrollment methods. You can configure it as a default setting:

1. In SmartConsole, click **Menu > Global properties > Advanced > Configure**.
2. Click **Certificates and PKI properties**.
3. Select **use_dkm_cert_by_default**.
4. Click **OK**.
5. Install the Access Control Policy.

Note - Generating certificates for Edge devices does not support DKM and will be generated locally on the management even if `use_dkm_cert_by_default` is enabled.

Configuration of PKI Operations

Trusting a CA - Step-By-Step

This section describes the procedures for obtaining a CA's own certificate, which is a prerequisite for trusting certificates issued by a CA.

In order to trust a CA, a CA server object has to be defined. The following sections deal with the various configuration steps required in different scenarios.

Trusting an ICA

A VPN module automatically trusts the ICA of the Security Management Server that manages it. No further configuration is required.

Trusting an Externally Managed CA

An externally managed CA is the ICA of another Security Management Server. The CA certificate has to be supplied and saved to disk in advance.

To establish trust:

1. In **Object Explorer**, click **New > Server > More > Trusted CA**.

The **Certificate Authority Properties** window opens.

2. Enter a **Name** for the CA object.
3. Go to the **OPSEC PKI** tab.
4. Click **Get**.
5. Browse to where you saved the peer CA certificate and select it.

The certificate details are shown.

Make sure the certificate's details are correct.

Make sure the SHA-1 and MD5 fingerprints of the CA certificate are correct.

6. Click **OK**.

Trusting an OPSEC Certified CA

The CA certificate has to be supplied and saved to the disk in advance.

Note - In case of SCEP automatic enrollment, you can skip this stage and fetch the CA certificate automatically after configuring the SCEP parameters.

The CA's Certificate must be retrieved either by downloading it with the CA options in the Certificate Authority object, or by obtaining the CA's certificate from the peer administrator in advance.

Define the CA object according to the following steps

1. In **Object Explorer**, click **New > Server > More > Trusted CA or Subordinate CA**.

The **Certificate Authority Properties** window opens.

2. Enter a **Name** for the CA object.
3. On the **OPSEC PKI** tab:

- For automatic enrollment, select **Automatically enroll certificate**.
- From the **Connect to CA with protocol**, select the protocol used to connect with the certificate authority, either SCEP, CMPV1 or CMPV2.

Note - For entrust 5.0 and later, use CMPV1.

4. Click **Properties**:

- If you chose SCEP as the protocol, in the **Properties for SCEP protocol** window, enter the CA identifier (such as example.com) and the Certification Authority/Registration Authority URL.
- If you chose CMPV1 as the protocol, in the **Properties for CMP protocol - V1** window, enter the applicable IP address and port number. (The default port is 829).
- If you chose CMPV2 as the protocol, in the **Properties for CMP protocol -V2** window, decide whether to use direct TCP or HTTP as the transport layer.

Note - If Automatic enrollment is not selected, then enrollment will have to be performed manually.

5. Choose a method for retrieving CRLs from this CA.

If the CA publishes CRLs on HTTP server choose **HTTP Server(s)**.

Certificates issued by the CA must contain the CRL location in an URL in the **CRL Distribution Point** extension.

If the CA publishes CRL on LDAP server, choose **LDAP Server(s)**.

In this case, you must define an LDAP Account Unit as well. See the [R81.10 Security Management Administration Guide](#) for more details about defining an LDAP object.

In the **LDAP Account Unit Properties** window, on the **General** tab, make sure to check the **CRL retrieval**.

Certificates issued by the CA must contain the LDAP DN on which the CRL resides in the CRL distribution point extension.

6. Click **Get**.

7. If SCEP is configured, it will try to connect to the CA and retrieve the certificate. If not, browse to where you saved the peer CA certificate and select it.

The certificate is fetched. Verify the certificate's details. Display and validate the SHA-1 and MD5 fingerprints of the CA certificate.

8. Click **OK**.

Certificate Revocation (All CA Types)

A certificate issued by the Internal Certificate Authority is revoked when the certificate object is removed. Otherwise, the CA administrator controls certificate revocation with the options on the **Advanced** tab of the CA object. In addition, the certificate must be removed from the Security Gateway.

A certificate cannot be removed if the Security Management Server infers from other settings that the certificate is in use, for example, that the Security Gateway belongs to one or more VPN communities and this is the only certificate of the Security Gateway.

To remove the certificate

1. Open the **IPsec VPN** tab of the applicable Security Gateway.
2. In the Repository of **Certificates Available to the Gateway**, select the applicable certificate and click **Remove**.

Certificate Recovery and Renewal

When a certificate is revoked or becomes expired, it is necessary to create another one or to refresh the existing one.

Recovery and Renewal with Internal CA

Removal of a compromised or expired certificate automatically triggers creation of a new certificate, with no intervention required by the administrator. To manually renew a certificate use the **Renew** button on the VPN page of the Security Gateway object.

Note - A Security Gateway can have only one certificate signed by one CA. When the new certificate is issued, you will be asked to replace the existing certificate signed by the same CA.

CA Certificate Rollover

CA Certificate Rollover is a VPN feature that enables rolling over the CA certificates used to sign client and Security Gateway certificates for VPN traffic, without risk of losing functionality at transition.

To achieve a gradual CA certificate rollover, CA Certificate Rollover enables VPN support for multiple CA certificates generated by third-party OPSEC-compliant CAs, such as Microsoft Windows CA. With multiple CA certificates, you can gradually rollover client and Security Gateway certificates during a transitional period when client and Security Gateway certificates signed by both CA certificates are recognized.

When a certificate is added to a CA that already has a certificate, the new certificate is defined as **Additional** and receives an index number higher by one than the highest existing certificate index number. The original certificate is defined as **Main**.

Only additional certificates can be removed. CA Certificate Rollover provides tools for adding and removing certificates, and for changing the status of a certificate from additional to main and from main to additional.

CA Certificate Rollover is for rolling over CA certificates with different keys. To add a CA certificate with the same key as the existing CA certificate (for example, to extend its expiration date), just Get the certificate from the OPSEC PKI tab of the CA properties, and do not use CA Certificate Rollover.

Managing a CA Certificate Rollover

With multiple CA certificates, you can gradually rollover client and Security Gateway certificates during a transitional period when client and Security Gateway certificates signed by both CA certificates are recognized.

This section describes a recommended workflow for the most common scenario. For full details of the CLI commands, see the **"CA Certificate Rollover CLI"** section.

Before you begin:

In SmartConsole, define a third-party OPSEC-compliant CA, such as Microsoft Windows CA, that is capable of generating multiple CA certificates. Generate the main CA certificate and define it in SmartConsole.

To roll over with a new CA certificate

1. Generate from the third-party CA a second CA certificate in DER format (PEM is not supported), with different keys than the previous CA certificate. Copy the certificate to the Security Management Server.
2. Connect to the command line on the Security Management Server.
3. Log in to the Expert mode.
4. Add the new CA certificate to the Security Management Server database's definitions for the third-party CA:

```
mcc add <CA Name> <Certificate File>
```

See ["mcc add" on page 265](#):

- *<CA Name>* is the name of the CA as defined in the Security Management Server database.
- *<Certificate File>* is the certificate filename or path.

5. The new CA certificate should now be defined as additional #1.

Make sure with the ["mcc lca" on page 268](#) or ["mcc show" on page 270](#) command.

6. In SmartConsole, install the Access Control Policy on all Security Gateways.

Use the new additional CA certificate to sign client certificates.

For performance reasons, as long as most clients still use certificates signed by the old CA certificate, you should leave the new CA certificate as the additional one and the old certificate as the main one. As long as the new CA certificate is not the main CA certificate, do not use it to sign any Security Gateway certificates.

CA Certificate Rollover CLI

To perform CA Certificate Rollover use the VPN Multi-Certificate CA commands - ["mcc" on page 263](#).

Adding Matching Criteria to the Validation Process

While certificates of an externally managed VPN entity are not handled by the local Security Management Server, you can still configure a peer to present a particular certificate when creating a VPN tunnel

Configuration

1. Open the **VPN** page of the externally managed VPN entity.
2. Click **Matching Criteria**.
3. Choose the characteristics of the certificate the peer is expected to present, including:
 - The CA that issued it
 - The exact DN of the certificate
 - The IP address that appears in the **Subject Alternate Name** extension of the certificate. (This IP address is compared to the IP address of the VPN peer itself as it appears to the VPN module during the IKE negotiation.)
 - The e-mail address appearing in the **Subject Alternate Name** extension of the certificate

CRL Cache Usage

To cancel or modify the behavior of the CRL Cache

1. Open the **Advanced Tab** of the Certificate Authority object.
2. To enable the CRL cache, check **Cache CRL on the Security Gateway**.

The cache should not be disabled for the ICA. In general, it is recommended that the cache be enabled for all CA types. The cache should be disabled (for non-ICAs) only if stringent security requirements mandate continual retrieval of the CRL.

Note - The ICA requires the use of a CRL cache, and should never be disabled.

3. If CRL Cache is enabled, choose whether a CRL is deleted from the cache when it expires or after a fixed period of time (unless it expires first). The second option encourages retrieval of a CRL more often as CRLs may be issued more frequently than the expiry time. By default a CRL is deleted from the cache after 24 hours.

Modifying the CRL Pre-Fetch Cache

To modify the duration of the Pre-fetch cache

1. In SmartConsole, click **Menu > Global Properties > Advanced > Configure**.
2. Click **Certificates and PKI properties**.
3. In the **prefetch_crls_duration** field, configure the duration.
4. Click **OK**.
5. Install the Access Control Policy.

Configuring CRL Grace Period

To configure the CRL Grace Period values

1. In SmartConsole, click **Menu > Global properties > VPN > Advanced**.
2. In the **CRL Grace Period** section, configure the applicable times.
3. Click **OK**.
4. Install the Access Control Policy.

The Grace Period can be defined for both the periods before and after the specified CRL validity period.

Configuring OCSP

To use OCSP, you must configure the CA object to use the OCSP revocation information method instead of the CRL method.

Use Database Tool (GuiDBEdit Tool) (see [sk13009](#)) to change the value of the field **ocsp_validation** to **true**. When set to true, the CA uses OCSP to make sure that certificates are valid. This is configured on the root CA and is inherited by the subordinate CAs.

To configure a CA to use OCSP, in Database Tool (GuiDBEdit Tool):


See [sk37803](#) for detailed instructions.

Domain Based VPN

Overview of Domain-based VPN

Domain Based VPN controls how VPN traffic is routed between Security Gateways within a community. To route traffic to a host behind a Security Gateway, you must first define the VPN domain for that Security Gateway. Configuration for VPN routing is done with SmartConsole or in the VPN routing configuration files on the Security Gateways.

In this figure, one of the host machines behind Security Gateway A tries to connect to a host computer behind Security Gateway B. For technical or policy reasons, Security Gateway A cannot establish a VPN tunnel with Security Gateway B. With VPN Routing, Security Gateways A and B can establish VPN tunnels through Security Gateway C.



Item	Description
A	Security Gateway A
B	Security Gateway B
C	Security Gateway C

VPN Routing and Access Control

VPN routing connections are subject to the same access control rules as any other connection. If VPN routing is correctly configured but a Security Policy rule exists that does not allow the connection, the connection is dropped. For example: a Security Gateway has a rule which forbids all FTP traffic from inside the internal network to anywhere outside. When a peer Security Gateway opens an FTP connection with this Security Gateway, the connection is dropped.

For VPN routing to succeed, a single rule in the Security Policy Rule Base must cover traffic in both directions, inbound and outbound, and on the central Security Gateway. To configure this rule, see ["Domain Based VPN" on the previous page](#).

Configuring VPN Routing in Domain Based VPN

Configure most common VPN routing scenarios through a VPN star community in SmartConsole.

You can also configure VPN routing between Security Gateways in the Security Management Server configuration file `$FWDIR/conf/vpn_route.conf`.

You can only configure VPN routing between Security Gateways that belong to a VPN community.

Configuring VPN Routing for Security Gateways in SmartConsole

To configure a VPN Routing in a star community in SmartConsole:

1. On the **Star Community** window, in the:
 - a. **Center Gateways** section, select the Security Gateway that functions as the "Hub".
 - b. **Satellite Gateways** section, select Security Gateways as the "spokes", or satellites.
2. On the **VPN Routing** page, **Enable VPN routing for satellites** section, select one of these options:
 - **To center and to other Satellites through center** - This allows connectivity between the Security Gateways, for example if the spoke Security Gateways have dynamically assigned IP addresses, and the Hub is a Security Gateway with a static IP address.

- **To center, or through the center to other satellites, to internet and other VPN targets** - This allows connectivity between the Security Gateways as well as the ability to inspect all communication passing through the Hub to the Internet.
3. Create an applicable Access Control Policy rule. Remember: one rule must cover traffic in both directions.
 4. NAT the satellite Security Gateways on the Hub if the Hub is used to route connections from Satellites to the Internet.

Configuration in the VPN Configuration File

For more granular control over VPN routing, edit the `$FWDIR/conf/vpn_route.conf` file on the Security Management Server.

The configuration file, `$FWDIR/conf/vpn_route.conf`, is a text file that contains the name of network objects.

The format is: **Destination, Next hop, Install on Security Gateway** (with tabbed spaces separating the elements).

Consider a simple VPN routing scenario consisting of Center gateway (hub) and two Satellite gateways (spokes). All machines are controlled from the same Security Management Server, and all the Security Gateways are members of the same VPN community. Only Telnet and FTP services are to be encrypted between the Satellites and routed through the Center:

Although you can do this easily in a VPN Star community, you can achieve the same goal if you edit the `$FWDIR/conf/vpn_route.conf` file:

Destination	Next hop router interface	Install on
Spoke_B_VPN_Dom	Hub_C	Spoke_A
Spoke_A_VPN_Dom	Hub_C	Spoke_B

In this instance, `Spoke_B_VPN_Dom` is the name of the network object group that contains spoke B's VPN domain. Hub C is the name of the Security Gateway enabled for VPN routing. `Spoke_A_VPN_Dom` is the name of the network object that represents Spoke A's encryption domain.

Example of the file contents:

```
Spoke_B_VPN_DOM Hub_C Spoke_A
Spoke_A_VPN_DOM Hub_C Spoke_B
```


Configuring the 'Accept VPN Traffic Rule'

In SmartConsole:

1. Double click on a Star or Meshed Community.
2. On the **Encrypted Traffic** page, select **Accept all encrypted traffic**.
3. In a Star community, choose between accepting encrypted traffic on **Both center and satellite gateways** or **Satellite gateways only**.
4. Click OK.

Configuring Multiple Hubs

Consider two Hubs, A and B. Hub A has two spokes, spoke_A1, and spoke_A2. Hub B has a single spoke, spoke_B. In addition, Hub A is managed from Security Management Server A, while Hub B is managed via Security Management Server B:



For the two VPN star communities, based around Hubs A and B:

- Spokes A1 and A2 need to route all traffic going outside of the VPN community through Hub A
- Spokes A1 and A2 also need to route all traffic to one another through Hub A, the center of their star community
- Spoke B needs to route all traffic outside of its star community through Hub B

A_community is the VPN community of A plus the spokes belonging to A. B_community is the VPN community. Hubs_community is the VPN community of Hub_A and Hub_B.

Configuring VPN Routing and Access Control on Security Management Server A

The `$FWDIR/conf/vpn_route.conf` file on Security Management Server 1 looks like this:

Destination	Next hop router interface	Install on
Spoke_B_VPN_Dom	Hub_A	A_Spokes

Destination	Next hop router interface	Install on
Spoke_A1_VPN_Dom	Hub_A	Spoke_A2
Spoke_A2_VPN_Dom	Hub_A	Spoke_A1
Spoke_B_VPN_Dom	Hub_B	Hub_A

Spokes A1 and A2 are combined into the network group object "A_spokes".

The applicable rule in the Security Policy Rule Base looks like this:

Source	Destination	VPN	Service	Action
*Any	*Any	A_Community B_Community Hubs_Community	*Any	Accept

Configuring VPN Routing and Access Control on Security Management Server B

The `$FWDIR/conf/vpn_route.conf` file on Security Management Server 2 looks like this:

Destination	Next hop router interface	Install On
Spoke_A1_VPN_Dom	Hub_B	Spoke_B
Spoke_A2_VPN_Dom	Hub_B	Spoke_B
Spoke_A1_VPN_Dom	Hub_A	Hub_B
Spoke_A2_VPN_Dom	Hub_A	Hub_B

The applicable rule in the Security Policy Rule Base looks like this:

Source	Destination	VPN	Service	Action
*Any	*Any	B_Community A_Community Hubs_Community	*Any	Accept

For both `$FWDIR/conf/vpn_route.conf` files:

- "A_Community" is a star VPN community comprised of Hub_A, Spoke_A1, and Spoke_A2
- "B_Community" is a star VPN community comprised of Hub_B and Spoke_B

- "Hubs-Community" is a *meshed* VPN community comprised of Hub_A and Hub_B (it could also be a star community with the central Security Gateways meshed).



VPN with One or More LSM Profiles

You can configure a VPN star community between two SmartLSM Profiles. The procedures below show a SmartLSM Gateway Profile and SmartLSM Cluster Profile. You can also configure the community with two SmartLSM Cluster Profiles or two SmartLSM Gateway Profiles. All included SmartLSM Gateway and SmartLSM Cluster Profiles must have the IPsec VPN blade enabled.

The procedure requires configuration in:

- SmartConsole
- CLI on the Security Management Server
- SmartProvisioning GUI
- CLI on the Center Security Gateway

Route Based VPN

Overview of Route-based VPN

The use of VPN Tunnel Interfaces (VTI) is based on the idea that setting up a VTI between peer Security Gateways is similar to connecting them directly.

A VTI is a virtual interface that can be used as a Security Gateway to the VPN domain of the peer Security Gateway. Each VTI is associated with a single tunnel to a Security Gateway. The tunnel itself with all of its properties is defined, as before, by a VPN Community linking the two Security Gateways. Configure the peer Security Gateway with a corresponding VTI. The native IP routing mechanism on each Security Gateway can then direct traffic into the tunnel as it would for other interfaces.


All traffic destined to the VPN domain of a peer Security Gateway is routed through the "associated" VTI. This infrastructure allows dynamic routing protocols to use VTIs. A dynamic routing protocol daemon running on the Security Gateway can exchange routing information with a neighboring routing daemon running on the other end of an IPsec tunnel, which appears to be a single hop away.

Route Based VPN can only be implemented between Security Gateways within the same VPN community.

To deploy Route Based VPN, Directional Rules have to be configured in the Rule Base of the Security Management Server. See ["Directional Enforcement within a Community" on page 169](#)

VPN Tunnel Interface (VTI)

A VPN Tunnel Interface is a virtual interface on a Security Gateway that is related to a VPN tunnel and connects to a remote peer. You create a VTI on each Security Gateway that connects to the VTI on a remote peer. Traffic routed from the local Security Gateway via the VTI is transferred encrypted to the associated peer Security Gateway.



In this scenario:

- There is a VTI connecting "Cluster GWA" and "GWb" (you must configure the same Tunnel ID on these peers)
- There is a VTI connecting "Cluster GWA" and "GWc" (you must configure the same Tunnel ID on these peers)
- There is a VTI connecting "GWb" and "GWc" (you must configure the same Tunnel ID on these peers)

A virtual interface behaves like a point-to-point interface directly connected to the remote peer. Traffic between network hosts is routed into the VPN tunnel with the IP routing mechanism of the Operating System. Security Gateway objects are still required, as well as VPN communities (and access control policies) to define which tunnels are available. However, VPN encryption domains for each peer Security Gateway are no longer necessary. The decision whether or not to encrypt depends on whether the traffic is routed through a virtual interface. The routing changes dynamically if a dynamic routing protocol (OSPF/BGP) is available on the network.

When a connection that originates on GWb is routed through a VTI to GWc (or servers behind GWc) and is accepted by the implied rules, the connection leaves GWb in the clear with the local IP address of the VTI as the source IP address. If this IP address is not routable, return packets will be lost.

The solution for this issue is:

- Configure a static route on GWb that redirects packets destined to GWc from being routed through the VTI

- Not including it in any published route
- Adding route maps that filter out GWc's IP addresses

Having excluded those IP addresses from route-based VPN, it is still possible to have other connections encrypted to those addresses (i.e. when not passing on implied rules) by using domain based VPN definitions.

The VTI can be configured in two ways:

VTI Type	Description
Numbered	You configure a local and remote IP address for each numbered VPN Tunnel Interface (VTI). For each Security Gateway, you configure a local IP address, a remote address, and the local IP address source for outbound connections to the tunnel. The remote IP address must be the local IP address on the remote peer Security Gateway. More than one VTI can use the same IP Address, but they cannot use an existing physical interface IP address.
Unnumbered	For unnumbered VTIs, you define a proxy interface for each Security Gateway. Each Security Gateway uses the proxy interface IP address as the source for outbound traffic. Unnumbered interfaces let you assign and manage one IP address for each interface. Proxy interfaces can be physical or loopback interfaces.

Using Dynamic Routing Protocols

VTIs allow the ability to use Dynamic Routing Protocols to exchange routing information between Security Gateways.

The Dynamic Routing Protocols supported on Gaia are:

- BGP4
- OSPFv2
- RIPv1
- RIPv2

VTIs in a Clustered Environment

When configuring numbered VTIs in a clustered environment, a number of issues need to be considered:

- Each member must have a unique source IP address.
- Every interface on each member requires a unique IP address.
- All VTIs going to the same remote peer must have the same name.
- Cluster IP addresses are required.

Configuring VTIs in Gaia Operating System

See the [R81.10 Gaia Administration Guide](#) > Chapter *Network Management* > Section *Network Interfaces* > Section *VPN Tunnel Interfaces*.

Note - For VTIs between Gaia Security Gateways and Cisco GRE gateways, you must manually configure the Hello/Dead packet intervals at 10/40 on the Gaia Security Gateways, or at 30/120 on the peer gateway. If not, OSPF is not able to get into the "FULL" state.

Enabling Route Based VPN


If you configure a Security Gateway for Domain Based VPN and Route Based VPN, Domain Based VPN takes precedence by default.

To force Route Based VPN to take priority, you must create a dummy (empty) group and assign it to the VPN domain.

To force Route-Based VPN to take priority:

1. In SmartConsole, from the left navigation panel, click **Gateways & Servers**.
2. Open the Security Gateway / Cluster object.
3. From the left tree, click **Network Management > VPN Domain**.
4. Select **User defined**.
5. Click the [...] button.
6. Click **New > Group > Simple Group**.
7. Enter a **Name**.
8. Click **OK** (leave this Group object empty).

Configuring Numbered VTIs - Example



The Security Gateways in this scenario are:

Device Type	Specific Computers
ClusterXL	Cluster GWa: - member_GWb1 - member_GWb2
VPN peers	- GWb - GWc

VTIs connect these Security Gateways:

- Members of "Cluster GWa" and "GWb"
- Members of "Cluster GWa" and "GWc"
- "GWb" and "GWc"

IP Configuration:

Peer	Type of IP Address and Interface	IP Address / Netmask
Cluster GWa	External Unique IP address of <code>eth0</code>	170.170.1.1 / 24
member_GWa1	External VIP address of <code>eth0</code>	170.170.1.10 / 24
	IP address of Sync interface <code>eth1</code>	5.5.5.1 / 24
	IP address of VTI for "GWb"	Local: 10.0.1.11 / 24 Remote: 10.0.0.2 / 24
	VIP address of VTI for "GWb"	10.0.1.10 / 24
	IP address of VTI for "GWc"	Local: 10.0.1.21 / 24 Remote: 10.0.0.3 / 24
	VIP address of VTI for "GWc"	10.0.1.20 / 24
Cluster GWa	External Unique IP address of <code>eth0</code>	170.170.1.2 / 24
member_GWa2	External VIP address of <code>eth0</code>	170.170.1.10 / 24
	IP address of Sync interface <code>eth1</code>	5.5.5.2 / 24
	IP address of VTI for "GWb"	Local: 10.0.1.12 / 24 Remote: 10.0.0.2 / 24
	VIP address of VTI for "GWb"	10.0.1.10 / 24
	IP address of VTI "vt-GWc"	Local: 10.0.1.22 / 24 Remote: 10.0.0.3 / 24
	VIP address of VTI for "GWc"	10.0.1.20 / 24
GWb	External Unique IP address of <code>eth0</code>	180.180.1.1 / 24
	IP address of VTI for "Cluster GWa"	Local: 10.0.0.2 / 24 Remote: 10.0.1.10 / 24
	IP address of VTI "vt-GWc"	Local: 10.0.0.2 / 24 Remote: 10.0.0.3 / 24
GWc	External Unique IP address of <code>eth0</code>	190.190.1.1 / 24
	IP address of VTI for "Cluster GWa"	Local: 10.0.0.3 / 24 Remote: 10.0.1.20 / 24

Peer	Type of IP Address and Interface	IP Address / Netmask
	IP address of VTI for "GWb"	Local: 10.0.0.3 / 24 Remote: 10.0.0.2 / 24

The example configurations below use the same Security Gateway names and IP addresses that are described in Numbered VTIs.

1. Configure the required VTIs on 'member_GWa1'

See the [R81.10 Gaia Administration Guide](#) > Chapter *Network Management* > Section *Network Interfaces* > Section *VPN Tunnel Interfaces*.

a. Configure a Numbered VPN Tunnel Interface for **GWb**.

Use these settings for the VTI:

Parameter	Value
VPN Tunnel ID	Integer from 1 to 99  Important - You must configure the same ID for GWb on all Cluster Members.
Peer	GWb
VPN Tunnel Type	Numbered
Local Address	10.0.1.11
Remote Address	10.0.0.2

- b. Configure a Numbered VPN Tunnel Interface for **GWc**.

Use these settings for the VTI:

Parameter	Value
VPN Tunnel ID	Integer from 1 to 99  Important - You must configure the same ID for GWc on all Cluster Members.
Peer	GWc
VPN Tunnel Type	Numbered
Local Address	10.0.1.21
Remote Address	10.0.0.3

2. Configure the required VTIs on 'member_GWa2'

See the [R81.10 Gaia Administration Guide](#) > Chapter *Network Management* > Section *Network Interfaces* > Section *VPN Tunnel Interfaces*.

- a. Configure a Numbered VPN Tunnel Interface for **GWb**.

Use these settings for the VTI:

Parameter	Value
VPN Tunnel ID	Integer from 1 to 99  Important - You must configure the same ID for GWb on all Cluster Members.
Peer	GWb
VPN Tunnel Type	Numbered
Local Address	10.0.1.12
Remote Address	10.0.0.2

- b. Configure a Numbered VPN Tunnel Interface for **GWc**.

Use these settings for the VTI:

Parameter	Value
VPN Tunnel ID	Integer from 1 to 99  Important - You must configure the same ID for GWc on all Cluster Members.
Peer	GWc
VPN Tunnel Type	Numbered
Local Address	10.0.1.22
Remote Address	10.0.0.3

3. Configure the required VTIs on 'GWb'

See the [R81.10 Gaia Administration Guide](#) > Chapter *Network Management* > Section *Network Interfaces* > Section *VPN Tunnel Interfaces*.

- a. Configure a Numbered VPN Tunnel Interface for **Cluster GWa**.

Use these settings for the VTI:

Parameter	Value
VPN Tunnel ID	Integer from 1 to 99  Important - You must configure the same ID you configured on all Cluster Members for GWb .
Peer	ClusterGWa
VPN Tunnel Type	Numbered
Local Address	10.0.0.2
Remote Address	10.0.1.10

- b. Configure a Numbered VPN Tunnel Interface for **GWc**.

Use these settings for the VTI:

Parameter	Value
VPN Tunnel ID	Integer from 1 to 99  Important - You must configure the same ID for this VTI on GWb and GWc .
Peer	GWc
VPN Tunnel Type	Numbered
Local Address	10.0.0.2
Remote Address	10.0.0.3

4. Configure the required VTIs on 'GWc'

See the [R81.10 Gaia Administration Guide](#) > Chapter *Network Management* > Section *Network Interfaces* > Section *VPN Tunnel Interfaces*.

- a. Configure a Numbered VPN Tunnel Interface for **Cluster GWa**.

Use these settings for the VTI:

Parameter	Value
VPN Tunnel ID	Integer from 1 to 99  Important - You must configure the same ID you configured on all Cluster Members for GWc .
Peer	ClusterGWa
VPN Tunnel Type	Numbered
Local Address	10.0.0.3
Remote Address	10.0.1.20

- b. Configure a Numbered VPN Tunnel Interface for **GWb**.

Use these settings for the VTI:

Parameter	Value
VPN Tunnel ID	Integer from 1 to 99  Important - You must configure the same ID for this VTI on GWc and GWb .
Peer	GWb
VPN Tunnel Type	Numbered
Local Address	10.0.0.3
Remote Address	10.0.0.2

5. Configure the Cluster object in SmartConsole

After configuring the VTIs on the cluster members, you must configure the Cluster Virtual IP addresses of these VTIs in the cluster object in SmartConsole.

- a. From the left navigation panel, click **Gateways & Servers**.
- b. Right-click the cluster object and select **Edit**.
- c. From the left tree, click **Network Management**.
- d. Click **Get Interfaces > Get Interfaces Without Topology**.

The VTIs appear in the **Topology** column as **Point to point**.

Interfaces are members of the same VTI if these criteria match:


- Peer
 - Remote IP address
 - Interface name
- e. Configure the Cluster Virtual IP addresses on the VTIs:
 - i. Select the VTI interface and click **Edit**.
 - ii. On the **General** page, enter the Virtual IP address.
 - iii. Click **OK**.

Virtual IP Addresses:

Name	Topology	Virtual IP	member_ GWa1	member_ GWa2	Comment
vpnt1	Point to point	10.0.1.10	10.0.1.11	10.0.1.12	VTI with GWb
vpnt2	Point to point	10.0.1.20	10.0.1.21	10.0.1.22	VTI with GWc

- f. Click **OK**.
- g. Install the Access Control Policy on the cluster object.

Enabling Dynamic Routing Protocols on VTIs - Example



The example below shows how the OSPF dynamic routing protocol is enabled on VTIs.

Note that the network commands for single members and cluster members are not the same.

For more information on VTIs and advanced routing commands, see the:

- [R81.10 Gaia Administration Guide](#).
- [R81.10 Gaia Advanced Routing Administration Guide](#).

When peering with a Cisco GRE enabled device, a point to point GRE tunnel is required.

Configuration:

OSPF configuration on 'member_GWa1'

vpnt1 is the VTI between 'member_GWa1' and 'GWb'

vpnt2 is the VTI between 'member_GWa1' and 'GWc'

```
member_GWa1:0> set ospf area 0.0.0.0 on
member_GWa1:0> set router-id 170.170.1.10
member_GWa1:0> set ospf interface vpnt1 area 0.0.0.0 on
member_GWa1:0> set ospf interface vpnt2 area 0.0.0.0 on
member_GWa1:0> set route-redistribution to ospf2 from kernel all-ipv4-routes on
member_GWa1:0> save config
member_GWa1:0> show configuration ospf
```

OSPF configuration on 'member_GWa2'

vpnt1 is the VTI between 'member_GWa2' and 'GWb'

vpnt2 is the VTI between 'member_GWa2' and 'GWc'

```
member_GWa2:0> set ospf area 0.0.0.0 on
member_GWa2:0> set router-id 170.170.1.10
member_GWa2:0> set ospf interface vpnt1 area 0.0.0.0 on
member_GWa2:0> set ospf interface vpnt2 area 0.0.0.0 on
member_GWa2:0> set route-redistribution to ospf2 from kernel all-ipv4-routes on
member_GWa2:0> save config
member_GWa2:0> show configuration ospf
```

OSPF configuration on 'GWb'

vpnt1 is the VTI between 'GWb' and 'Cluster GWa'

vpnt3 is the VTI between 'GWb' and 'GWc'

```
GWb:0> set ospf area 0.0.0.0 on
GWb:0> set router-id 180.180.1.1
GWb:0> set ospf interface vpnt1 area 0.0.0.0 on
GWb:0> set ospf interface vpnt3 area 0.0.0.0 on
GWb:0> set route-redistribution to ospf2 from kernel all-ipv4-routes on
GWb:0> save config
GWb:0> show configuration ospf
```

OSPF configuration on 'GWc'

vpnt2 is the VTI between 'GWc' and 'Cluster GWa'

vpnt3 is the VTI between 'GWc' and 'GWb'

```
GWc:0> set ospf area 0.0.0.0 on
GWc:0> set router-id 190.190.1.1
GWc:0> set ospf interface vpnt2 area 0.0.0.0 on
GWc:0> set ospf interface vpnt3 area 0.0.0.0 on
GWc:0> set route-redistribution to ospf2 from kernel all-ipv4-routes on
GWc:0> save config
GWc:0> show configuration ospf
```

Configuring Anti-Spoofing on VTIs in SmartConsole

1. From the left navigation panel, click **Gateways & Servers**.
2. Right-click the Security Gateway object and select **Edit**.
3. From the left tree, click **Network Management**.
4. Select a VTI interface, and click **Edit**.
5. From the left tree, click **General**.
6. In the **Topology** section, click **Modify**.
7. In the **IP Addresses behind peer Security Gateway that are within reach of this interface** section, select:
 - **Not Defined** - To accept all traffic.
 - **Specific** - To choose a particular network. The IP addresses in this network will be the only addresses accepted by this interface.
8. In the **Perform Anti-Spoofing based on interface topology** section, select **Don't check packets from** to make sure Anti-Spoofing does not occur for traffic from IP addresses from certain internal networks to the external interface. Configure a **Network** object that represents those internal networks with valid addresses, and from the drop-down list, select that Network object.

Anti-Spoofing does not apply to objects selected in the **Don't check packets from** drop-down menu.
9. In the **Spoof Tracking** field, select the applicable options.
10. Click **OK**.
11. Install the Access Control Policy on the Security Gateway object.

Routing Multicast Packets Through VPN Tunnels

Multicast is used to transmit a single message to a select group of recipients. IP Multicasting applications send one copy of each datagram (IP packet) and address it to a group of computers that want to receive it. This technique addresses datagrams to a group of receivers (at the multicast address) rather than to a single receiver (at a unicast address). The network is responsible for forwarding the datagrams to only those networks that need to receive them. PIM is required for this feature.

For more about Multicasting, see the [R81.10 Security Management Administration Guide](#) > Chapter *Creating an Access Control Policy* > Section *Multicast Access Control*.

Multicast traffic can be encrypted and forwarded across VPN tunnels that were configured with VPN tunnel interfaces (virtual interfaces associated with the same physical interface). All participant Security Gateways, both on the sending and receiving ends, must have a virtual interface for each VPN tunnel and a multicast routing protocol must be enabled on all participant Security Gateways.

To enable multicast service on a Security Gateway functioning as a rendezvous point, add a rule to the security policy of that Security Gateway to allow only the specific multicast service to be accepted unencrypted, and to accept all other services only through the community. Corresponding Access Control rules enabling multicast protocols and services should be created on all participating Security Gateways.

For example:

Source	Destination	VPN	Services & Applications	Action	Track
Multicast Security Gateways	Multicast Security Gateways	Any	igmp pim	Accept	Log
Sample Host	Multicast Group Address	Sample Community	Multicast Service Group	Accept	Log

Large Scale VPN

A VPN that connects branch offices, worldwide partners, remote clients, and other environments, can reach hundreds or thousands of peers. A VPN on this scale brings new challenges.

Each time a new VPN peer is deployed in production configuration and policy installation is required for all participating VPN Gateways.

Large Scale VPN (LSV) addresses these challenges and facilitates deployment without the need for peer configuration and policy installation.

Configuring LSV

Workflow:

1. Configure the Certificate Authority.
2. Configure the Center VPN Security Gateway.
3. Configure the VPN community.
4. Configure the LSV profile.
5. Install the Security Policy.

Configuring LSV

This configuration is applied on the central VPN Gateway

1. Configure the Certificate Authority.

The CA certificate has to be supplied and saved to the disk in advance.

Note - In case of SCEP automatic enrollment, you can skip this stage and fetch the CA certificate automatically after configuring the SCEP parameters.

The CA's Certificate must be retrieved either by downloading it with the CA options in the Certificate Authority object, or by obtaining the CA's certificate from the peer administrator in advance.

Define the CA object according to the following steps

- a. In **Object Explorer**, click **New > Server > More > Trusted CA** or **Subordinate CA**.

The **Certificate Authority Properties** window opens.

- b. Enter a **Name** for the CA object.
- c. On the **OPSEC PKI** tab:
 - For automatic enrollment, select **Automatically enroll certificate**.
 - From the **Connect to CA with protocol**, select the protocol used to connect with the certificate authority, either SCEP, CMPV1 or CMPV2.

Note - For entrust 5.0 and later, use CMPV1.
- d. Click **Properties**:
 - If you chose SCEP as the protocol, in the **Properties for SCEP protocol** window, enter the CA identifier (such as example.com) and the Certification Authority/Registration Authority URL.
 - If you chose CMPV1 as the protocol, in the **Properties for CMP protocol - V1** window, enter the applicable IP address and port number. (The default port is 829).
 - If you chose CMPV2 as the protocol, in the **Properties for CMP protocol - V2** window, decide whether to use direct TCP or HTTP as the transport layer.

Note - If Automatic enrollment is not selected, then enrollment will have to be performed manually.
- e. Choose a method for retrieving CRLs from this CA.

If the CA publishes CRLs on HTTP server choose **HTTP Server(s)**.

Certificates issued by the CA must contain the CRL location in an URL in the **CRL Distribution Point** extension.

If the CA publishes CRL on LDAP server, choose **LDAP Server(s)**.

In this case, you must define an LDAP Account Unit as well. See the [R81.10 Security Management Administration Guide](#) for more details about defining an LDAP object.

In the **LDAP Account Unit Properties** window, on the **General** tab, make sure to check the **CRL retrieval**.

Certificates issued by the CA must contain the LDAP DN on which the CRL resides in the CRL distribution point extension.
- f. Click **Get**.

- g. If SCEP is configured, it will try to connect to the CA and retrieve the certificate. If not, browse to where you saved the peer CA certificate and select it.

The certificate is fetched. Verify the certificate's details. Display and validate the SHA-1 and MD5 fingerprints of the CA certificate.

- h. Click **OK**.

2. Configure the certificate for the central VPN Security Gateway.

The devices participating in the LSV community must all share a signed certificate from the same Certificate Authority signed for the Central VPN Gateway.

A certificate is automatically issued by the Internal Certificate Authority for all internally managed entities that are VPN-capable. That is, after the administrator enables the **IPsec VPN Software Blade** in a Security Gateway or Cluster object (on the **General Properties** page > on the **Network Security** tab).

The process for obtaining a certificate from an **OPSEC PKI CA** or **External Check Point CA** is identical.

Manual Enrollment with OPSEC Certified PKI

To create a PKCS#10 Certificate Request:

- a. Create a Certificate Authority object.
- b. From the left navigation panel, click **Gateways & Servers**.
- c. Double-click the applicable Security Gateway or Cluster object.
- d. From the left tree, click **General Properties** and make sure to enable the **IPsec VPN Software Blade**.
- e. From the left tree, click **IPsec VPN**.
- f. In the section **Repository of Certificates Available to the Gateway**, click **Add**.

The **Certificate Properties** window opens.

- g. In the **Certificate Nickname** field, enter a text string.

The nickname is only an identifier and has no bearing on the content of the certificate.

- h. From the drop-down menu **CA to enroll from**, select the Certificate Authority that issues the certificate.

Note - The menu shows only trusted Certificate Authorities and subordinate Certificate Authorities that lead directly to a trusted Certificate Authority. If the CA that issues the certificate is a subordinate CA that does not lead directly to a trusted CA, it is not in the menu.

- i. In the section **Key pair generation and storage**, select the applicable method:

- **Store keys on the Security Management server** - Certificate creation is performed entirely between the Management Server and applicable CA. The keys and the certificate are downloaded securely to the Security Gateway (Cluster Members) during policy installation.
- **Store keys on the Module** - Management Server directs the Security Gateway (or Cluster Members) to create the keys and supply only the required material for creation of the certificate request. Only the certificate is downloaded to the Security Gateway (Cluster Members) during policy installation.

- j. Click **Generate**.

The **Generate Certificate Properties** window opens.

- k. Enter the applicable DN.

The CA administrator determines the final DN that appears in the certificate.

If a **Subject Alternate Name** extension is required in the certificate, select **Define Alternate Name**.

The public key and the DN are then used to DER-encode a PKCS#10 Certificate Request.



Note - Adding the object's IP address as the **Alternate Name** extension can be configured as a default setting.

This configuration also applies for Internal Certificate Authorities.

- i. In SmartConsole, click **Menu > Global properties > Advanced > Configure**.
 - ii. Click **Certificates and PKI properties**.
 - iii. Select these options:
 - **add_ip_alt_name_for_ICA_certs** (closer to the top of this page)
 - **add_ip_alt_name_for_opsec_certs** (closer to the bottom of this page)
 - iv. Click **OK** to close the **Advanced Configuration** window.
 - v. Click **OK** to close the **Global properties** window.
- I. When the certificate appears in the section **Repository of Certificates Available to the Gateway**:
 - i. Select this certificate.
 - ii. Click **View**.
 - iii. In the **Certificate View** window:
 - i. Click inside the window.
 - ii. Select the whole text (press the **CTRL+A** keys, or right-click the mouse and click **Select All**).
 - iii. Copy the whole text (press the **CTRL+C** keys, or right-click the mouse and click **Copy**).
 - iv. Paste the text into a plain text editor (like Notepad).
 - v. Click **OK**.

- m. Send the certificate information to the Certificate Authority administrator.

The CA administrator must now complete the task of issuing the certificate.

Different CAs provide different ways of doing this, such as an advanced enrollment form (as opposed to the regular form for users).

The issued certificate may be delivered in various ways, for example, email.

- n. After the certificate arrives from the Certificate Authority administrator, you must save it in the Certificate Authority object:
 - i. In SmartConsole, click **Objects > Object Explorer** (or press the **CTRL+E** keys).
 - ii. In the left tree, click **Servers**.
 - iii. Double-click the applicable Certificate Authority object.
 - iv. Click the **OPEC PKI** tab.
 - v. In the **Certificate** section, click **Get**.
 - vi. Browse to the location, where you saved the certificate file.
 - vii. Select the certificate file and click **Open**.
 - viii. If the certificate details are correct, click **OK** to accept this certificate.
 - ix. Click **OK** to close the **Certificate Authority Properties** window.
 - x. Close the **Object Explorer** window.
- o. Publish the SmartConsole session

Automatic Enrollment with the Certificate Authority

On the **OPSEC PKI** tab of the Certificate Authority object:

- a. Select the option **Automatically enroll certificate**.
- b. Select the applicable protocol - **scep** or **cmp**.

Follow these steps:

- a. From the left navigation panel, click **Gateways & Servers**.
- b. Double-click the applicable Security Gateway or Cluster object.
- c. From the left tree, click **General Properties** and make sure to enable the **IPsec VPN Software Blade**.
- d. From the left tree, click **IPsec VPN**.

- e. In the section **Repository of Certificates Available to the Gateway**, click **Add**.

The **Certificate Properties** window opens.

- f. In the **Certificate Nickname** field, enter a text string.

The nickname is only an identifier and has no bearing on the content of the certificate.

- g. From the drop-down menu **CA to enroll from**, select the Certificate Authority that issues the certificate.

Note - The menu shows only trusted CAs and subordinate CAs that lead directly to a trusted CA. If the CA that issues the certificate is a subordinate CA that does not lead directly to a trusted CA, it is not in the menu.

- h. In the section **Key pair generation and storage**, select the applicable method:

- **Store keys on the Security Management server** - Certificate creation is performed entirely between the Management Server and applicable CA. The keys and the certificate are downloaded securely to the Security Gateway (Cluster Members) during policy installation.
- **Store keys on the Module** - Management Server directs the Security Gateway (or Cluster Members) to create the keys and supply only the required material for creation of the certificate request. Only the certificate is downloaded to the Security Gateway (Cluster Members) during policy installation.

- i. Click **Generate** and select **Automatic enrollment**.

The **Generate Keys and Get Automatic Enrollment Certificate** window opens.

- Supply the **Key Identifier** and your secret **Authorization code**.
- Click **OK**.

- j. When the certificate appears in the section **Repository of Certificates Available to the Gateway**:
 - i. Select this certificate.
 - ii. Click **View**.
 - iii. In the **Certificate View** window, click **Copy to Clipboard** or **Save to File**.
- k. Send the request to CA administrator.

Different Certificate Authorities provide different means for doing this. For example, an advanced enrollment form on their website. The issued certificate can be delivered in various ways, such as by email. After you receive the certificate, save it to disk.
- l. From the left tree click, **IPsec VPN**.
- m. In the section **Repository of Certificates Available to the Gateway**:
 - i. Select the applicable certificate.
 - ii. Click **Complete**.
- n. Browse to the folder where you stored the issued certificate, select the certificate, and examine the certificate details.
- o. Click **OK** to close the Security Gateway or Cluster object.
- p. Publish the SmartConsole session

Enrolling through a Subordinate CA

When enrolling through a Subordinate CA:

- Supply the password of the Subordinate CA which issues the certificate (not the CA at the top of the hierarchy).
- The Subordinate CA must lead directly to a trusted CA.

3. Configure the VPN community.

Configuring a new VPN community

- a. From the left navigation panel, click **Security Policies**.
- b. In the top left section **Access Control**, click **Policy**.
- c. In the bottom left section **Access Tools**, click **VPN Communities**.
- d. Click **New** (✱) and select **Star Community**.


- e. Enter a name for the VPN Community.
 - f. In the **Center Gateways** area, click the + icon to add one or more Security Gateways (Clusters) to be in the center of the community.
 - g. In the **Satellite Gateways** area, click the + icon to add one or more Security Gateways (Clusters) to be around the center Security Gateways (Clusters).
 - h. Click **OK**.
- The Community uses the default encryption and VPN Routing settings.
- i. **Optional:** Edit more settings for the VPN Community in the community object.

4. Configure the LSV profile.

Configuring the LSV Profile

- a. Edit the VPN Community object.
- b. From the left tree, click **Gateways**.
- c. In the **Satellite Gateways** section, click the + icon > **New (*)** > **Large Scale VPN**.

The **New Large Scale VPN Profile** window opens.



- d. In the **Certificate Authority** section, select the applicable CA object.
- e. **Optional:** Configure the VPN domain for the LSV profile

You can limit the number of IP addresses used in an encryption domain of each satellite VPN Gateway and restrict the VPN access to specific group of networks.

Important - If the Encryption Domain of the LSV gateways overlaps (the same or partial Encryption Domain is configured for two or more peer devices), the default behavior is to use the VPN connection of the peer the connected last. The kernel parameter "lsv_prefer_new_peer" on Security Gateways (Cluster Members) controls this behavior. The default value of this kernel parameter is 1.

5. Install the Security Policy.

Monitoring LSV Peers and Tunnels

To monitor the LSV Peers

You can monitor LSV peers on a Security Gateway with the **vpn lsv** command.

1. Connect to the command line on the Security Gateway (each Cluster Member).
2. Log in to the Expert mode.
3. Run:

```
vpn lsv
```

Output:

```
Select an option.
*****          Select Option          *****
(1)          List all LSV peers
(2)          Show LSV peer's details
(3)          Remove an LSV peer
(4)          Remove all LSV peers
(Q)          Quit
*****
```

Tunnel Management

Overview of Tunnel Management

The VPN tunnel transports data securely. You can manage the types of tunnels and the number of tunnels with these features:

- **Permanent Tunnels** - Keeps VPN tunnels active to allow real-time monitoring capabilities.
- **VPN Tunnel Sharing** - Provides greater interoperability and scalability between Security Gateways. It also controls the number of VPN tunnels created between peer Security Gateways.

See the status of all VPN tunnels in SmartView Monitor. For details see *Monitoring Tunnels* in the [R81.10 Logging and Monitoring Administration Guide](#).

Permanent Tunnels

As companies have become more dependent on VPNs for communication to other sites, uninterrupted connectivity has become more crucial than ever before. Therefore it is essential to make sure that the VPN tunnels are kept up and running. Permanent Tunnels are constantly kept active and as a result, make it easier to recognize malfunctions and connectivity problems. Administrators can monitor the two sides of a VPN tunnel and identify problems without delay.

Each VPN tunnel in the community may be set to be a Permanent Tunnel. Since Permanent Tunnels are constantly monitored, if the VPN tunnel is down, then a log, alert, or user defined action, can be issued. A VPN tunnel is monitored by periodically sending "tunnel test" packets. As long as responses to the packets are received the VPN tunnel is considered "up." If no response is received within a given time period, the VPN tunnel is considered "down." Permanent Tunnels can only be established between Check Point Security Gateways. The configuration of Permanent Tunnels takes place on the community level and:

- Can be specified for an entire community. This option sets every VPN tunnel in the community as permanent.
- Can be specified for a specific Security Gateway. Use this option to configure specific Security Gateways to have permanent tunnels.
- Can be specified for a single VPN tunnel. This feature allows configuring specific tunnels between specific Security Gateways as permanent.

Permanent Tunnels in a MEP Environment

In a *Multiple Entry Point* (MEP) environment, VPN tunnels that are active are rerouted from the predefined primary Security Gateway to the backup Security Gateway if the primary Security Gateway becomes unavailable. When a Permanent Tunnel is configured between Security Gateways in a MEP environment where RIM is enabled, the satellite Security Gateways see the center Security Gateways as "unified." As a result, the connection will not fail but will fail over to another center Security Gateway on a newly created permanent tunnel. For more information on MEP see ["Multiple Entry Point \(MEP\) VPNs" on page 174](#).

Tunnel Testing for Permanent Tunnels

Check Point uses a proprietary protocol to test if VPN tunnels are active, and supports any site-to-site VPN configuration. Tunnel testing requires two Security Gateways, and uses UDP port 18234. Check Point tunnel testing protocol does not support 3rd party Security Gateways.

Terminating Permanent Tunnels

Once a Permanent Tunnel is no longer required, the tunnel can be shut down. Permanent Tunnels are shut down by deselecting the configuration options to make them active and re-installing the policy.

Dead Peer Detection

In addition to Tunnel Testing, *Dead Peer Detection* (DPD) is a different method to test if VPN tunnels are active. *Dead Peer Detection* does support 3rd party Security Gateways and supports permanent tunnels with interoperable devices based on IKEv1/IKEv2 DPD (IKEv1 DPD is based on RFC 3706). It uses IPsec traffic patterns to minimize the number of messages required to confirm the availability of a peer.

The tunnel testing mechanism is the recommended keepalive mechanism for Check Point to Check Point VPN gateways because it is based on IPsec traffic and requires an IPsec established tunnel. DPD is based on IKE encryption keys only.

Permanent Tunnel Based on Dead Peer Detection

DPD can monitor remote peers with the permanent tunnel feature. All related behavior and configurations of permanent tunnels are supported.

 **Important** - When configuring DPD with a third-party VPN gateway, you must include the Check Point Security Gateway's external interface IP address in the encryption domain of the third-party VPN peer.

To configure DPD for a permanent tunnel, the permanent tunnel must be in the VPN community. After you configure the permanent tunnel, configure Permanent Tunnel mode Based on DPD. There are different possibilities for permanent tunnel mode:

- **tunnel_test** (default) - The permanent tunnel is monitored by a tunnel test (as in earlier versions). It works only between Check Point Security Gateways. Keepalive packets are always sent.
- **dpd** - The active DPD mode. A peer receives DPD requests at regular intervals (10 seconds). DPD requests are only sent when there is no traffic from the peer.
- **passive** - The passive DPD mode. Peers do not send DPD requests to this peer. Tunnels with passive peers are monitored only if there is IPsec traffic and incoming DPD requests.

Note: To use this mode for only some gateways, enable the `forceSendDPDPayload` registry key on Check Point remote peers.

To enable DPD monitoring:

On each VPN gateway in the VPN community, configure the **tunnel_keepalive_method** property, in [Database Tool \(GuiDBEdit Tool\)](#) or dbedit (see [sk13301](#)). This includes 3rd Party gateways. (You cannot configure different monitor mechanisms for the same gateway).

1. In Database Tool (GuiDBEdit Tool), go to **Network Objects > network_objects > <Name of Security Gateways object> > VPN**.
2. For the **Value**, select a permanent tunnel mode.
3. Save all the changes.
4. Install the Access Control Policy.

Optional Configuration:

- **IKE Initiation Prevention** - By default, when a valid IKE SA is not available, a DPD request message triggers a new IKE negotiation. To prevent this behavior, set the property **dpd_allowed_to_init_ike** to **false**.

Edit the property in [Database Tool \(GuiDBEdit Tool\)](#) > **Network Objects > network_objects > <Name of Security Gateways object> > VPN**.

- **Delete IKE SAs for dead peer** - Based on RFC 3706, a VPN Gateway has to delete IKE SAs from a dead peer. This functionality is enabled, by default.
 - To disable the feature, add this line to the `$CPDIR/tmp/.CPprofile.sh` file and then **reboot**:

```
DPD_DONT_DEL_SA=0 ; export DPD_DONT_DEL_SA
```

Note - It is not supported to change the value of this environment variable in the current shell session with the "**export DPD_DONT_DEL_SA=0**" command.

- To enable the feature (if you disabled it), remove the line with "DPD_DONT_DEL_SA" from the `$CPDIR/tmp/.CPprofile.sh` file and then **reboot**.

Note - It is not supported to change the value of this environment variable in the current shell session with the `"export DPD_DONT_DEL_SA=1"` command.

VPN Tunnel Sharing

For a VPN community, the VPN tunnel sharing configuration is set on the **Tunnel Management** page of the **Community Properties** window.

For a specific Security Gateway, the configuration is set on the **VPN Advanced** page of the Security Gateway properties window.

Tunnel test is a proprietary Check Point protocol used to see if VPN tunnels are active. Tunnel testing requires two Security Gateways and uses UDP port 18234. Third party gateways do not support tunnel testing.

VPN Tunnel Sharing provides interoperability and scalability by controlling the number of VPN tunnels created between peer Security Gateways.

There are three available settings:

- **One VPN tunnel per each pair of hosts**
- **One VPN tunnel per subnet pair**
- **One VPN tunnel per Security Gateway pair**

In case of a conflict between the tunnel properties of a VPN community and a Security Gateway object that is a member of that same community, the "stricter" setting is followed. For example, a Security Gateway that was set to **One VPN Tunnel per each pair of hosts** and a community that was set to **One VPN Tunnel per subnet pair**, would follow **One VPN Tunnel per each pair of hosts**.

Configuring Tunnel Features

To configure Tunnel Management options:

1. In SmartConsole, click **Object Explorer** (Ctrl+E)
2. Click **New > VPN Community** and choose **Star Community** or **Meshed community**.
3. Click **Tunnel Management**. and configure the tunnel settings:

- Permanent Tunnels
- Tracking Options
- VPN Tunnel Sharing

Permanent Tunnels

In the **Star Community** or **Meshed community** object, on the **Tunnel Management** page, select **Set Permanent Tunnels**.

These are the options:

- **On all tunnels in the community**
- **On all tunnels of specific Security Gateways**
- **On specific tunnels in the community**

To configure all tunnels as permanent, select **On all tunnels in the community**. Clear this option to terminate all Permanent Tunnels in the community.

To configure on all tunnels of specific Security Gateways:

1. Select **On all tunnels of specific gateways** and click **Select Gateways**.

The **Select Gateway** window opens.

To terminate Permanent Tunnels connected to a specific Security Gateway, select the Security Gateway object and click **Remove**.

2. To configure the **Tracking** options for a specific Security Gateway, select a Security Gateway object and click **Gateway Tunnel Properties**.

To configure on specific tunnels in the community:

1. Select **On specific tunnels in the community** and click **Select Permanent Tunnels**.

The **Select Permanent Tunnels** window opens.

2. Double click in the white cell that intersects the Security Gateways where a permanent tunnel is required.

The **Tunnel Properties** window opens.

3. Click **Set these tunnels to be permanent tunnels**.

To terminate the Permanent Tunnel between these two Security Gateways, clear **Set these tunnels to be permanent tunnels**.

4. Click **OK**.

Advanced Permanent Tunnel Configuration

You can configure advanced VPN settings globally. In addition, you can configure DPD thresholds per community.

To configure advanced VPN settings globally:

1. In SmartConsole, click **Menu > Global properties**.
2. Click **Advanced > Configure**.
3. Click **VPN Advanced Properties > Tunnel Management** to see the attributes that may be configured to customize the amount of tunnel tests sent and the intervals in which they are sent:
 - **life_sign_timeout** - Set the amount of time the tunnel test or DPD runs without a response before the peer host is declared 'down.'
 - **life_sign_transmitter_interval** - Set the time between tunnel tests or DPD.
 - **life_sign_retransmissions_count** - When a tunnel test does not receive a reply, another test is resent to confirm that the peer is 'down.' The Life Sign Retransmission Count is set to how many times the tunnel test is resent without receiving a response.
 - **life_sign_retransmissions_interval** - Set the time between the tunnel tests that are resent after it does not receive a response from the peer.
 - **cluster_status_polling_interval** - (applicable for High Availability Clusters only) - Set the time between tunnel tests between a primary Security Gateway and a backup Security Gateway. The tunnel test is sent by the backup Security Gateway. When there is no reply, the backup Security Gateway will become active.
4. Click **OK**.
5. If you changed the existing setting, then install the Access Control Policy.

To define DPD thresholds per community:

You can define the DPD thresholds `life_sign_timeout` and `life_sign_transmitter` per community.

1. In Check Point Database Tool (GuiDBedit), in the left upper pane , select the **Tables** tab.
2. Select **Managed Objects**.
3. Select **communities**.
4. In the right upper pane, select a VPN community.

5. In the lower pane, scroll to the `life_sign_timeout` row and double click on the number in the **Value** column. Or, right click on the number in the **Value** column and click **Edit**.

A window opens..

6. Enter a number of seconds.
 - Minimum: 5 seconds
 - Maximum: 3,600 seconds
 - Default: 40 seconds.
7. Click **OK**.
8. Scroll to the `life_sign_transmitter_interval` row and double click on the number in the **Value** column. Or, right click on the number in the **Value** column and click **Edit**.

A window opens.

9. Enter a number of seconds.
 - Minimum: 5 seconds
 - Maximum: 600 seconds
 - Default: 10 seconds
10. Click **OK**.

Tracking Options

You can configure alerts to stay updated on the status of permanent VPN tunnels.

To configure logs and alerts for VPN tunnel status:

1. In the properties of the VPN Community, open the **Tunnel Management** page.
2. In **Tunnel down track**, select the alert when a tunnel is down.
3. In **Tunnel up track**, select the alert when a tunnel is up.

The alerts are configured for the tunnels that are defined as permanent, based on the settings on the page.

See status of all VPN tunnels in SmartView Monitor.

To open SmartView Monitor:

1. In SmartConsole, click **Logs & Monitor**.
2. Click **New Tab**.

3. From the bottom of this page, click **Tunnel & User Monitoring**.

For more details, see *Monitoring Tunnels* in the [R81.10 Logging and Monitoring Administration Guide](#).

Route Injection Mechanism

Overview of Route Injection

Route Injection Mechanism (RIM) enables a Security Gateway to use dynamic routing protocols to propagate the encryption domain of a VPN peer Security Gateway to the internal network. When a VPN tunnel is created, RIM updates the local routing table of the Security Gateway to include the encryption domain of the VPN peer.

Note - Route Injection is not currently supported for IPv6.

RIM can only be enabled when permanent tunnels are configured for the community. Permanent tunnels are kept alive by tunnel test packets. When a Security Gateway fails to reply, the tunnel is considered "down." As a result, RIM deletes the route to the failed link from the local routing table, which triggers neighboring dynamic routing enabled devices to update their routing information accordingly. This results in a redirection of all traffic destined to travel across the VPN tunnel, to a pre-defined alternative path.


There are two possible methods to configure RIM:

- Automatic RIM - RIM automatically injects the route to the encryption domain of the peer Security Gateways.
- Custom Script - Specify tasks for RIM to perform according to specific needs.

Route injection can be integrated with MEP functionality, which sends return packets back through the same MEP Security Gateway. For more information on MEP, see ["Multiple Entry Point \(MEP\) VPNs" on page 174](#).

Automatic RIM

In this scenario:



Label	Meaning
R1	Router 1
G1	Security Gateway 1
R2	Router 2
G2	Security Gateway 2

- RIM is enabled in the VPN community in which Security Gateway 1 and Security Gateway 2 participate.
- When the Security Gateways create a VPN tunnel, and the Permanent Tunnel status changes to "UP", RIM updates the routing tables:
 - The routing table on Security Gateway 1 gets the routes for the encryption domain of Security Gateway 2.
 - The routing table on Security Gateway 2 gets the routes for the encryption domain of Security Gateway 1.
- Security Gateway 1 has a dynamic routing neighborhood with Router 1 and propagates RIM routes to Router 1.
- Security Gateway 2 has a dynamic routing neighborhood with Router 2 and propagates RIM routes to Router 2.

- If the VPN tunnel becomes unavailable (Permanent Tunnel status with a peer changes to "DOWN"), then RIM removes routes from the routing tables of Security Gateway 1 and Security Gateway 2. The Security Gateways update their corresponding neighbors: Router 1 and Router 2. The routers start to send traffic over the leased line.

Below are the routing tables on the Security Gateways and Routers based on the diagram above. Entries in bold represent routes that RIM injected into the Security Gateway's local routing table:

For Security Gateway 1:

Destination	Netmask	Next Hop	Metric
0.0.0.0	0.0.0.0	172.16.10.2	1
192.168.21.0	255.255.255.0	172.16.10.2	1
192.168.11.0	255.255.255.0	192.168.10.1	1

Security Gateway 2:

Destination	Netmask	Next Hop	Metric
0.0.0.0	0.0.0.0	172.16.20.2	1
192.168.11.0	255.255.255.0	172.16.20.2	1
192.168.21.0	255.255.255.0	192.168.20.1	1

Router 1 (behind Security Gateway 1):

Destination	Netmask	Next Hop	Metric
0.0.0.0	0.0.0.0	192.168.10.2	1
192.168.21.0	255.255.255.0	192.168.10.2	1
192.168.21.0	255.255.255.0	10.10.10.2	2

Router 2 (behind Security Gateway 2):

Destination	Netmask	Next Hop	Metric
0.0.0.0	0.0.0.0	192.168.20.2	1

Destination	Netmask	Next Hop	Metric
192.168.11.0	255.255.255.0	192.168.20.2	1
192.168.11.0	255.255.255.0	10.10.10.1	2

Custom Scripts

Custom scripts can be run on any Security Gateway in the community. These scripts are executed whenever a tunnel changes its state (example: goes "up" or "down"). Such an event, for example, can be the trigger that initiates a dial-up connection.

There is a script template **custom_rim** (with a **.sh** or **.bat** extension depending on the operating system) in the **\$FWDIR/scripts/** directory.

Sample customized script:

```
#!/bin/sh

# This script is invoked each time a tunnel is configured with the
# RIM option
# and the tunnel changed state.
#
# You may add your custom commands to be invoked here.

# Parameters read from command line.
RIM_PEER_GW=$1
RIM_NEW_STATE=$2
RIM_HA_STATE=$3
RIM_FIRST_TIME=$4
RIM_PEER_ENC_NET=$5

case "${RIM_NEW_STATE}" in
    up)
        # Place your action for tunnels that came up
        ;;
    down)
        # Place your action for tunnel that went down
        ;;
esac
```

Where:

- **RIM_PEER_GW:** Peer Security Gateway
- **RIM_NEW_STATE:** Change in the state of the Security Gateway,(example: up or down).


- `RIM_HA_STATE`: State of a single Security Gateway in a cluster (example: standby or active).
- `RIM_FIRST_TIME`: The script is executed separately for each network in the peer's encryption domain. Though the script can be executed multiple times on a peer, this parameter is transferred to the script with the value of '1' only the first time the script runs on the peer. The value '1' indicates the first time this script is executed. The next time the script is executed, it is transferred with the value of '0' and the parameter is disregarded. For example, you can send an email alert to the system administrator the moment a tunnel goes down.
- `RIM_PEER_ENC_NET`: VPN domain of the VPN peer.

Injecting Peer Security Gateway Interfaces

You can inject the IP addresses of the peer Security Gateway into the routing tables, in addition to the networks behind the Security Gateway.

For example, after a VPN tunnel is created, RIM injects the encryption domain of the peer Security Gateway into the local routing tables of both Security Gateways. When RIM enabled Security Gateways communicate with a Security Gateway that has Hide NAT enabled, it is necessary to inject the peer's interfaces.

In this scenario:



- Security Gateways A and B are both RIM enabled and Security Gateway C has Hide NAT enabled on the external interface ("hiding" all the IP addresses behind it).
- Host 1, behind Security Gateway C, initiates a VPN tunnel with Host 2, through Security Gateway A.
- Router 3 holds routes to all the hosts behind Security Gateway C. Because Router 3 does not have the Hide NAT IP address of Security Gateway C, Router 3 cannot properly route packets back to Host 1.

To route back packets:

1. In SmartConsole:
 - a. Click **Menu > Global properties**.
 - b. Click **VPN Advanced Properties > Tunnel Management**.
 - c. Select **RIM_inject_peer_interfaces**. This injects Router 3 with all of the IP addresses of Security Gateway C (this includes the Hide NAT address).
 - d. Click **OK**.
 - e. Install the Access Control Policy.
2. Configure the router not to propagate the information injected to other Security Gateways. For example, in the scenario shown above this could result in Security Gateway B routing traffic to Security Gateway C through Security Gateway A.

Configuring RIM

Configuring RIM in a Star Community

1. In SmartConsole, click **Objects > Object Explorer** (or press **Ctrl E**).
2. From the left tree, select **VPN Communities**.
3. Open the applicable **Star Community** object.
4. From the left tree, click **Tunnel Management**.
5. In the **Permanent Tunnels** section, select **Set Permanent Tunnels**.

These *"Permanent Tunnels" on page 146* modes are available:

- On all tunnels in the community
- On all tunnels of specific Security Gateways
- On specific tunnels in the community

 **Note** - RIM can only be enabled on permanent tunnels. If Multiple Entry Point (MEP) is enabled on the community, you must select **On all tunnels in the community**. See ["Configuring Tunnel Features" on page 149](#).

6. Select **Enable Route Injection Mechanism (RIM)**.
7. Click **Settings**.

The **Star Community Settings** window opens.

In the **Community** section:

- **Enable automatic Route Injection Mechanism** - RIM runs automatically on the central or satellite Security Gateways.
- **Enable customer editable script execution** - A customized script runs on central or satellite Security Gateways whenever a tunnel changes its states (goes up or down).

In the **Tracking** section, configure the applicable tracking options: **Log**, **Popup Alert**, **Mail Alert**, **SNMP Trap Alert**, **User Defined Alert**

8. Click **OK** to close all configuration windows.
9. Close the Object Explorer.
10. Install the Access Control Policy.
11. If you selected **Enable customer editable script execution**, then you must edit the `$FWDIR/scripts/custom_rim.sh` script on each of the Security Gateways.

Configuring RIM in a Meshed Community

1. In SmartConsole, click **Objects** menu > **Object Explorer** (or press Ctrl E).
2. From the left tree, select **VPN Communities**.
3. Open the applicable **Meshed Community** object.
4. From the left tree, click **Tunnel Management**.
5. In the **Permanent Tunnels** section, select **Set Permanent Tunnels**.

These ["Permanent Tunnels" on page 146](#) modes are available:

- **On all tunnels in the community**
- **On all tunnels of specific Security Gateways**
- **On specific tunnels in the community**

 **Note** - RIM can only be enabled on permanent tunnels. If Multiple Entry Point (MEP) is enabled on the community, you must select **On all tunnels in the community**. See ["Configuring Tunnel Features" on page 149](#).

6. Select **Enable Route Injection Mechanism (RIM)**.
7. Click **Settings**.

The **Meshed Community Settings** window opens.

In the **Community** section:

- **Enable automatic Route Injection Mechanism** - RIM runs automatically on the central or satellite Security Gateways.
- **Enable customer editable script execution** - A customized script runs on central or satellite Security Gateways whenever a tunnel changes its state (goes up or down).

In the **Tracking** section:

- Configure the applicable tracking options:

Log, Popup Alert, Mail Alert, SNMP Trap Alert, User Defined Alert

8. Click **OK** to close all configuration windows.
9. Close the Object Explorer.
10. Install the Access Control Policy.
11. If you selected **Enable customer editable script execution**, then you must edit the `$FWDIR/scripts/custom_rim.sh` script on each of the Security Gateways.

Enabling the RIM_inject_peer_interfaces flag

To enable the **RIM_inject_peer_interfaces** flag:

1. In SmartConsole, click **Menu > Global properties**.
2. Click **Advanced > Configure**.
3. Click **VPN Advanced Properties > Tunnel Management**.
4. Select **RIM_inject_peer_interfaces**.
5. Click **OK**.
6. Install the Access Control Policy.

Configuring RIM in Gaia

RIM automatically configures routes for peer encryption domains in the Gaia OS kernel on Security Gateways:

- When a VPN tunnel is up, RIM adds the applicable routes.
- When a VPN tunnel is down, RIM removes the applicable routes.

If the VPN tunnel state changes from "DOWN" to "UP" again, it can take time for the routes for peer encryption domains to appear again in the routing table on Security Gateways. You can configure the Gaia OS kernel to keep these routes even when the VPN tunnel is down.

Configure the applicable settings on Security Gateways in Gaia Portal or in Gaia Clish.

- Configuration in Gaia Portal:
 1. In the tree view, click **Advanced Routing > Routing Options**.
 2. In the **Kernel Options** area, select the **Kernel Routes** option.
 3. Click **Apply**.
- Configuration in Gaia Clish:
 1. `set kernel-routes on`
 2. `save config`

Gaia Gateways in a Star VPN Community

For Gaia Security Gateways to publish routes, run these CLI commands on all Security Gateways at the center of the community.

For more information, see the [R81.10 Gaia Advanced Routing Administration Guide](#).

1. `set routemap <Routemap Name> id <ID Number>`

For example:

```
set routemap RIM id 5
```

2. `set routemap <Routemap Name> id <ID Number> match protocol kernel`

For example:

```
set routemap RIM id 5 match protocol kernel
```

3. `set ospf export-routemap <Routemap Name> preference 1 on`

For example:

```
set ospf export-routemap RIM preference 1 on
```

4. `set routemap <Routemap Name> id <ID Number> allow`

For example:

```
set routemap RIM id 5 allow
```

5. `set routemap <Routemap Name> id <ID Number> on`

For example:

```
set routemap RIM2 id 10 on
```

6. `set routemap <Routemap Name> id <ID Number> match nexthop <IP Address of OSPF Interface of the other RIM GW> on`

For example:

```
set routemap RIM2 id 10 match nexthop <10.16.50.3> on
```

7. `set routemap <Routemap Name> id <ID Number> restrict`

For example:

```
set routemap RIM2 id 10 restrict
```

8. `set ospf import-routemap <Routemap Name> preference 1 on`

For example:

```
set ospf import-routemap RIM2 preference 1 on
```

9. `save config`

Wire Mode

Overview of Wire Mode

The *Wire Mode* improves connectivity by allowing existing connections to fail over successfully by bypassing firewall enforcement. Traffic within a VPN community is, by definition, private and secure. In many cases, the firewall and the rule on the firewall concerning VPN connections is unnecessary. With the *Wire Mode*, the firewall can be bypassed for VPN connections by defining internal interfaces and communities as "trusted".

When a packet reaches a Security Gateway, the Security Gateway asks itself two questions regarding the packet(s):

Is this information coming from a "trusted" source?

Is this information going to a "trusted" destination?


If the answer to both questions is yes, and the VPN Community to which both Security Gateways belong is designated as "*Wire Mode* enabled," stateful inspection is not enforced and the traffic between the trusted interfaces bypasses the firewall. Since no stateful inspection takes place, no packets can be discarded. The VPN connection is no different from any other connection along a dedicated wire. This is the meaning of "*Wire Mode*." Since stateful inspection no longer takes place, dynamic routing protocols (which do not survive state verification in non-wire mode configuration) can now be deployed. *Wire Mode* thus facilitates ["Route Based VPN" on page 119](#).

Wire Mode Scenarios

Wire mode can be used to improve connectivity and performance in different infrastructures. This section describes scenarios that benefit from the implementation of wire mode.

Wire Mode in a MEP Configuration

In this scenario:




- Security Gateway **M1** and Security Gateway **M2** are both wire mode enabled and have trusted internal interfaces.
- The community where Security Gateway **M1** and Security Gateway **M2** reside, is wire mode enabled.
- Host 1, residing behind Security Gateway **S1** is communicating through a VPN tunnel with **Host 2** residing behind Security Gateway **M1**.
- MEP is configured for Security Gateway **M1** and Security Gateway **M2** with Security Gateway **M1** being the primary Security Gateway and Security Gateway **M2** as the backup.

In this case, if Security Gateway **M1** goes down, the connection fails over to Security Gateway **M2**. A packet leaving **Host 2** will be redirected by the router behind Security Gateway **M1** to Security Gateway **M2** since Security Gateway **M2** is designated as the backup Security Gateway. Without wire mode, stateful inspection is enforced at Security Gateway **M2** and the connection is dropped. Packets that come into a Security Gateway whose session was initiated through a different Security Gateway, are considered "out-of-state" packets. Since Security Gateway **M2**'s internal interface is "trusted," and wire mode is enabled on the community, no stateful inspection is performed and Security Gateway **M2** will successfully continue the connection without losing any information.

Wire Mode with Route Based VPN

In this scenario:




- Wire mode is enabled on Center Security Gateway C (without an internal trusted interface specified).
- The community is wire mode enabled.
- Host 1 residing behind Satellite Security Gateway A wishes to open a connection through a VPN tunnel with Host 2 behind Satellite Security Gateway B.

In a satellite community, Center Security Gateways are used to route traffic between Satellite Security Gateways within the community.

In this case, traffic from the Satellite Security Gateways is only rerouted by Security Gateway C and cannot pass through Security Gateway C's firewall. Therefore, stateful inspection does not need to take place at Security Gateway C. Since wire mode is enabled on the community and on Security Gateway C, making them trusted, stateful inspection is bypassed. Stateful inspection, however, does take place on Security Gateways A and B.

Wire Mode Between Two VPN Communities

In this scenario:



- Security Gateway A belongs to Community 1.
- Security Gateway B belongs to Community 2.
- Security Gateway C belongs to Communities 1 and 2.
- Wire mode is enabled on Center Security Gateway C (without an internal trusted interface specified).
- Wire mode is enabled on both communities.
- Host 1 residing behind Satellite Security Gateway A wishes to open a connection through a VPN tunnel with Host 2 behind Satellite Security Gateway B.

Wire mode can also be enabled for routing VPN traffic between two Security Gateways which are not members of the same community. Security Gateway C is a member of both communities and therefore recognizes both communities as trusted. When host 1 behind Security Gateway A initiates a connection to host 2 behind Security Gateway B, Security Gateway C is used to route traffic between the two communities. Since the traffic is not actually entering Security Gateway C, there is no need for stateful inspection to take place at that Security Gateway. Stateful inspection, however, does take place on Security Gateways A and B.

Special Considerations for Wire Mode

Wire mode does not work with IPv6.

Configuring Wire Mode

Wire mode is configured in two places:

- Community Properties (Meshed or Star)
- Security Gateway Properties

Enabling Wire Mode on a VPN Community

1. In SmartConsole, click the **Objects** menu > **Object Explorer**.
2. From the left tree, select the **VPN Communities**.
3. Open the VPN Community object.
4. From the left tree, click **Wire Mode**.
5. Select **Allow uninspected encrypted traffic between Wire mode interfaces of the Community members**.
6. To enable Wire Mode Routing, select **Wire Mode Routing - Allow members to route uninspected encrypted traffic in VPN routing configurations**.
7. Click **OK**.
8. Install the Access Control Policy.

Enabling Wire Mode on a Specific Security Gateway

1. In SmartConsole, from the left navigation panel, click **Gateways & Servers**.
2. Open the applicable Security Gateway object.
3. From the left tree, click **IPsec VPN > VPN Advanced**.
4. In the **Wire mode** section:
 - a. Select **Support Wire Mode (and Wire mode routing - route uninspected encrypted traffic in VPN routing configurations)**.
 - b. Click **Add**.
 - c. Select the interfaces to be trusted by the selected Security Gateway.
 - d. Click **OK**.
 - e. Select **Log Wire mode traffic** to log the Wire Mode activity.
5. Click **OK**.
6. Install the Access Control Policy.

Directional VPN Enforcement

Overview of Directional VPN

When a VPN community is selected in the VPN column of the Security Policy Rule Base, the source and destination IP addresses can belong to any of the Security Gateways in the community. In other words, the traffic is bidirectional; any of the Security Gateways can be the source of a connection, any of the Security Gateways can be the destination endpoint. But what if the administrator (in line with the company's security policy) wished to enforce traffic in one direction only? Or to allow encrypted traffic to or from Security Gateways **not** included in the VPN community? To enable enforcement within VPN communities, VPN implements Directional VPN.


Directional VPN specifies where the source address must be, and where the destination address must be. In this way, enforcement can take place:

- Within a single VPN community
- Between VPN communities

Directional Enforcement within a Community

The example figure below shows a simple meshed VPN community called *MyIntranet*.

VPN traffic within the MyIntranet Mesh is bidirectional. Meaning, either of the Security Gateways (or the hosts behind the Security Gateways in the VPN domains) can be the source or destination address for a connection.



Source	Destination	VPN	Service	Action	Track
*Any	*Any	MyIntranet => MyIntranet MyIntranet =>internal_clear internal_clear => MyIntranet	telnet	Accept	Log
*Any	*Any	MyIntranet	telnet	Accept	Log

The match conditions are represented by a series of compound objects. The match conditions enforce traffic in the following directions:

- To and from the VPN Community via VPN routing (**MyIntranet => MyIntranet**)
- From the Community to the local VPN domains (**MyIntranet =>internal_clear**)
- From the local VPN domains to the VPN community (**internal_clear => MyIntranet**)

Configurable Objects in a Direction

The table below shows all the objects that can be configured in a direction, including three new objects created for Directional VPN:

Name of Object	Description
Remote Access	Remote Access VPN community
Site to Site VPN	Regular Star or Mesh VPN community
Any Traffic	Any traffic
All_GwToGw	All Site to Site VPN communities
All_Communities	All Site to Site and Remote Access VPN communities
External_clear	For traffic outside the VPN community
Internal_clear	For traffic between local domains within the VPN community

Note - Clear text connections originating from these objects are not subject to enforcement:

- **Any Traffic**
- **External_clear**
- **Internal_clear**

There is *no limit* to the number of VPN directions that you can configure in a single rule. In general, if you have many directional enforcements, consider replacing them with a standard bidirectional condition.

Directional Enforcement between Communities

VPN Directional Enforcement can take place between two VPN communities. In this case, one Security Gateway must be configured as a member of both communities and the enforcement point between them. Every other peer Security Gateway in both communities must have a route entry to the enforcement point Security Gateway in its `$FWDIR/conf/vpn_route.conf` file.

To add a route entry to the enforcement point Security Gateway:


On the management module of each Security Gateway in the community (except for the enforcement point Security Gateway), add an entry in the `$FWDIR/conf/vpn_route.conf` file:

Destination	Next hop router interface	Install on
<destination_community_obj>	<enforcement_point_gw>	<managed_FW_object>

These are the variables in the entry:

- `destination_community_obj` - a network object for the combined encryption domain of the community
- `enforcement_point_gw` - the Security Gateway that is a member of both communities and transfers the encrypted traffic between them
- `managed_FW_object` - all community members that are managed by the management module

In the example below, Washington is a Mesh community, and London is a VPN Star.



The directional VPN rule below must be configured for the enforcement point Security Gateway in the Access Control Policy Rule Base:

Source	Destination	VPN	Services & Applications	Action
*Any	*Any	Washington => London	*Any	Accept

The rule is applied to all VPN traffic that passes through the enforcement point Security Gateway between the Washington and London communities. If a connection is opened from a source in the Washington Mesh, and the destination is in the London Star, the connection is allowed. Otherwise, the connection is denied.

Note - The Directional Enforcement applies only to the first packet of a connection. If the connection is permitted, the following packets of this connection are also permitted, including the packets in the opposite direction.

Configuring Directional VPN Within a Community

To configure Directional VPN within a community:

1. In SmartConsole, click **Menu > Global properties > VPN > Advanced**.
2. Select **Enable VPN Directional Match** in VPN Column.
3. Click **OK**.
4. In SmartConsole, from the left navigation panel, click **Security Policies**.
5. In the **Access Control** section, click in the applicable rule.

6. In the VPN column of this rule, select **Directional Match Condition**.

The **New Directional Match Condition** window opens.

7. In the **Traffic reaching from** drop-down box, select the object for **Internal_clear** (the source).
8. In the **Traffic leaving to** drop-down box, select the applicable VPN community object (the destination).
9. Add another directional match, in which the applicable VPN community object is both the source and destination.

This allows traffic from the local domain to the community, and within the community.

10. Click **OK**.
11. Install the Access Control Policy.

Configuring Directional VPN Between Communities

To configure Directional VPN between communities:

1. In SmartConsole, click **Menu > Global properties > VPN > Advanced**.
2. Select **Enable VPN Directional Match** in VPN Column.
3. Click **OK**.
4. In SmartConsole, from the left navigation panel, click **Security Policies**.
5. In the **Access Control** section, click in the applicable rule.
6. In the VPN column of this rule, select **Directional Match Condition**.

The **New Directional Match Condition** window opens.

7. In the **Traffic reaching from** drop-down box, select the source of the connection.
8. In the **Traffic leaving to** drop-down box, select the destination of the connection
9. Click **OK**.
10. Install the Access Control Policy.

Multiple Entry Point (MEP) VPNs

Overview of MEP

Multiple Entry Point (MEP) is a feature that provides a High Availability and Load Sharing solution for VPN connections. The VPN module is installed on a Security Gateway that provides a single point of entry to the internal network. This Security Gateway makes the internal network "available" to peer Security Gateways. If the Security Gateway becomes unavailable, then peer Security Gateways cannot access the internal network.

A MEP environment has two or more Security Gateways that enable access to the same VPN domain. In a MEP environment, peer Security Gateways provide redundancy for uninterrupted access to the internal network.

VPN High Availability Using MEP or Clustering

Both MEP and Clustering are ways of achieving High Availability and Load Sharing.

However:

- Unlike the members of a ClusterXL Security Gateway Cluster, there is no physical restriction on the location of MEP Security Gateways. MEP Security Gateways can be geographically separated machines. In a cluster, the clustered Security Gateways need to be in the same location, directly connected via a *sync* interface.
- MEP Security Gateways can be managed by different Security Management Servers. Cluster members must be managed by the same Security Management Server.
- In a MEP configuration there is no "state synchronization" between the MEP Security Gateways. In a cluster, all of the Security Gateways hold the "state" of all the connections to the internal network. If one of the Security Gateways fails, the connection passes seamlessly over (performs *failover*) to another Security Gateway, and the connection continues. In a MEP configuration, if a Security Gateway fails, the current connection is lost and one of the backup Security Gateways picks up the *next* connection.
- In a MEP environment, the decision which Security Gateway to use is taken on the remote side; in a cluster, the decision is taken on the Security Gateway side.

Implementation

MEP is implemented using RDP for Check Point Security Gateways and DPD for 3rd party Gateways / Cloud vendors.

- RDP is a proprietary Probing Protocol (PP) that sends special UDP RDP packets to port 259 to discover whether an IP is reachable. This protocol is proprietary to Check Point and does not conform to RDP as specified in RFC 908 / RFC 1151.
Note - These UDP RDP packets are not encrypted, and only test the availability of a peer.
- DPD is a different method that discovers whether an IP is reachable. It supports 3rd party Security Gateways / Cloud vendors based on IKEv1/IKEv2.

 **Note** - In an MEP environment, a Security Gateway determines which protocol to use automatically.


The peer continuously probes or polls all MEP Security Gateways in order to discover which of the Security Gateways are "up", and chooses a Security Gateway according to the configured selection mechanism. Since RDP/DPD packets are constantly being sent, the status of all Security Gateways is known and updated when changes occur. As a result, all Security Gateways that are "up" are known.

There are two available methods to implement MEP:

MEP Method	Description
Explicit MEP	Only Star communities with more than one central Security Gateway can enable explicit MEP. This MEP method provides multiple entry points to the network behind the Security Gateways. When available, Explicit MEP is the recommended method.
Implicit MEP	This MEP method is supported in all scenarios, where fully or partially overlapping encryption domains exist, or where Primary-Backup Security Gateways are configured.

Explicit MEP

In a Site To Site Star VPN community, explicit MEP is configured in the VPN community object. When MEP is enabled, the satellites consider the "unified" VPN domain of all the Security Gateways as the VPN domain for each Security Gateway. This unified VPN domain is considered the VPN domain of each Security Gateway:



In the figure, a Star VPN community has two central Security Gateways, **M1** and **M2** (for which MEP has been enabled), and three satellite Security Gateways - **S1**, **S2**, and **S3**. When **S2** opens a connection with **Host-1** (which is behind **M1** and **M2**), the session is initiated through either **M1** or **M2**. Priority among the MEP Security Gateways is determined by the MEP entry point selection mechanism.

If **M2** is the selected entry point and becomes unavailable, the connection to **Host-1** fails over to **M1**. Returning packets will be rerouted with RIM or IP Pool NAT. For more information about returning packets, see the section "Routing Return Packets".

There are four methods used to choose which of the Security Gateways will be used as the entry point for any given connection:

Method	Description
Select the closest Security Gateway to source	First to respond
Select the closest Security Gateway to destination	By VPN domain
Random selection	For Load distribution
Manually set priority list	MEP rules

If you select either **By VPN domain** , or **Manually set priority list**, then **Advanced** options provide additional granularity.


MEP Selection Methods

MEP Selection Method	Description
First to Respond	The first Security Gateway to reply to the peer Security Gateway is chosen. An organization would choose this option if, for example, the organization has two Security Gateways in a MEP configuration - one in London, the other in New York. It makes sense for VPN peers located in England to try the London Security Gateway first and the NY Security Gateway second. Being geographically closer to the peers in England, the London Security Gateway will be the first to respond, and becomes the entry point to the internal network. See "Overview of the "First to Respond" method" below.
VPN Domain	If the destination IP address belongs to a particular VPN domain, the Security Gateway of that domain becomes the chosen entry point. This Security Gateway becomes the Primary Security Gateway, while other Security Gateways in the MEP configuration become its Backup Security Gateways. See "Overview of the "By VPN Domain" method" on the next page.
Random Selection	The remote peer randomly selects a Security Gateway, with which to open a VPN connection. For each source/destination IP address pair, a new Security Gateway is randomly selected. An organization might have a number of Security Gateways with equal performance abilities. In this case, it makes sense to enable load distribution to use these Security Gateways in a random and equal way. See "Overview of the "Random Selection" method" on page 180.
Manually set priority list	Priorities of Security Gateways can be set manually for the entire VPN community, or for individual satellite Security Gateways. See "Overview of the "Manually Set Priority List" method" on page 180.

Overview of the "First to Respond" method

When there is no primary Security Gateway, all Security Gateways share "equal priority".

When all Security Gateways share equal priority:



1. Remote peers send RDP/DPD packets to all the Security Gateways in the MEP configuration.
2. The first Security Gateway to respond to the probing RDP/DPD packets gets chosen as the entry point to network.


The idea behind *first to respond* is proximity. The Security Gateway, which is "closer" to the remote peer, responds first.

3. A VPN tunnel is opened with the first to respond. All subsequent connections pass through the chosen Security Gateway.
4. If the Security Gateway ceases to respond, a new Security Gateway is chosen.

Overview of the "By VPN Domain" method

Before you enable MEP, each IP address belongs to a specific VPN domain. With **By VPN Domain**, the Security Gateway of that domain becomes the chosen entry point.

In the example figure below, the VPN Star community has two central MEP Security Gateways (**M1** and **M2**, each with its own VPN domain), and remote satellite **S1**.



Host-2 (in the VPN domain of satellite **S1**) initiates a connection with **Host-1**. The connection can be directed through either **M1** or **M2**. However, **Host-1** is within **M2**'s original VPN domain. For this reason, **M2** is considered the Security Gateway "closest" to the destination IP address. **M2** is therefore considered the primary Security Gateway and **M1** the backup Security Gateway for **Host-1**. If there were additional Security Gateways in the center, these Security Gateways would also be considered as backup Security Gateways for **M2**.

If the VPN domains have fully or partially overlapping encryption domains, then more than one Security Gateway will be chosen as the "closest" entry point to the network. As a result, more than one Security Gateway will be considered as "primary." When there are more than one primary or backup Security Gateways available, the Security Gateway is selected with an additional selection mechanism. This advanced selection mechanism can be either (see the section "**Advanced Settings**"):

- **First to Respond**
- **Random Selection** (for load distribution)

For return packets you can use RIM on the center Security Gateways. If RIM is also enabled, set a metric with a lower priority value for the leased line than the VPN tunnel. The satellite **S1** might simultaneously have more than one VPN tunnel open with the MEP Security Gateways, for example **M2** as the chosen entry point for **Host-1** and **M1** as the chosen entry point for **Host-3**. While both **M1** and **M2** will publish routes to **Host-1** and **Host-3**, the lower priority metric will ensure the leased line is used only when one of the Security Gateways goes down.

Overview of the "Random Selection" method

With this method, a different Security Gateway is randomly selected as an entry point for incoming traffic. Evenly distributing the incoming traffic through all the available Security Gateways can help prevent one Security Gateway from becoming overwhelmed with too much incoming traffic.

The Security Gateways are probed with RDP/DPD packets, as in all other MEP configurations, to create a list of responding Security Gateways. A Security Gateway is randomly chosen from the list of responding Security Gateways. If a Security Gateway stops responding, another Security Gateway is (randomly) chosen.

A new Security Gateway is randomly selected for every source/destination IP addresses pair. While the source and destination IP addresses remain the same, the connection continues through the chosen Security Gateway.


In such a configuration, RIM is not supported. IP Pool NAT must be enabled to ensure return packets are correctly routed through the chosen Security Gateway.

Overview of the "Manually Set Priority List" method

The Security Gateway that will be chosen (from the central Security Gateways in the Star VPN community) as the entry point to the core network can be controlled by manually setting a priority per source Security Gateway.

Each priority constitutes a MEP Rule.

In the figure below, three MEP members (**M1**, **M2**, **M3**) provide entry points to the network for three satellite Security Gateways (**S1**, **S2**, **S3**). Satellite **S1** can be configured to try the Security Gateways in the following order: **M1**, **M2**, **M3**, giving the highest priority to **M1**, and the lowest priority to **M3**. Satellite **S2** can be configured to try the Security Gateways in the following order: **M2**, **M3** (but not to try **M1**).



Each of these priorities constitutes a MEP rule in the MEP manual priority list window:

Item	Description
1	Default MEP Rule
2	First MEP Rule
3	Second MEP Rule

The MEP manual priority list window is divided into the default rule, and rules which provide exceptions to the default rule. The default MEP rule takes effect when:

- No MEP rules are defined
- When the source of the connection cannot be found in the Exception priority rules

The Exception priority rules section contains three priority levels: primary, secondary, and tertiary. While there are only three priority levels,

- The same priority can be assigned to several central Security Gateways
- The same rule can be assigned to several satellite Security Gateways
- A priority level can be left blank

In the second MEP rule below:

Central Security Gateways **M3** and **M1** have equal priority. The same rule is being applied to satellites **S2** and **S3**.

When more than one Security Gateway is assigned the same priority level, which Security Gateway will be chosen is resolved according to the **Advanced** settings.

Advanced Settings

In some instances, more than one Security Gateway is available in the center with no obvious priority between them. For example - as shown in the second example of the second MEP rule, above - more than one Security Gateway is assigned "second" priority. In this scenario, **Advanced** options are used to decide which Security Gateway is chosen: **First to Respond**, or **Random Selection**. (Choose **Random Selection** to enable load balancing between the Security Gateways.)

When "manually set priority list" is the MEP selection mechanism, *RIM is supported*. RIM can be configured with "manually set priority list" because the "random selection" mechanism available on the **Advanced** button is different from the random selection mechanism used for MEP.

For the "random selection" mechanism employed for MEP, a different Security Gateway is selected for each source/destination IP addresses pair. For the random selection mechanism available from the **Advanced** button, a single MEP entry point is randomly selected and then used for all connections, and does not change according to source/destination pair. Load distribution is therefore achieved since every satellite Security Gateway is randomly assigned a Security Gateway as its entry point. This makes it possible to enable RIM at the same time.

Tracking

If the **Tracking** option is enabled for MEP, this information is logged by each satellite Security Gateway:

- The resolved peer Security Gateway (a Security Gateway in the MEP)
- The priority of the resolved Security Gateway (primary, secondary, tertiary)
- Whether the resolved Security Gateway is responding

For example, in the scenario shown in the section "**Manually Set Priority List**", satellite **S1** opens a connection to the VPN domain that includes Security Gateways **M1**, **M2**, and **M3**. **M1** is the resolved peer. If tracking is enabled, the log reads:

```
Resolved peer for tunnel from S1 to the MEP that contains M1, M2,
and M3, is: M1 (Primary Security Gateway, responding).
```

Implicit MEP

There are three methods to implement implicit MEP:

Method	Description
First to Respond	The first Security Gateway to reply to the peer Security Gateway is chosen. An organization would choose this option if, for example, the organization has two Security Gateways in a MEP configuration - one in London, the other in New York. It makes sense for VPN peers located in England to try the London Security Gateway first and the NY Security Gateway second. Being geographically closer to the peers in England, the London Security Gateway will be the first to respond, and becomes the entry point to the internal network.  Note - First to Respond MEP is configured by default. See "Overview of the "Implicit First to Respond" method" on the next page.
Primary-Backup	One or multiple backup Security Gateways provide "high availability" for a primary Security Gateway. The remote peer is configured to work with the primary Security Gateway, but switches to the backup Security Gateway if the primary goes down. An organization might decide to use this configuration if it has two Security Gateways in a MEP environment, one of which is stronger than the other. It makes sense to configure the stronger Security Gateway as the primary. Or perhaps both Security Gateways are the same in terms of strength of performance, but one has a cheaper or faster connection to the Internet. In this case, the Security Gateway with the better Internet connection should be configured as the primary. See "Overview of the "Implicit Primary-Backup Security Gateways" method" on page 185 and "Configuring the "Implicit Primary-Backup" method" on page 188.
Load Distribution	The remote peer randomly selects a Security Gateway, with which to open a VPN connection. For each source/destination IP address pair, a new Security Gateway is randomly selected. An organization might have a number of Security Gateways with equal performance abilities. In this case, it makes sense to enable load distribution to use these Security Gateways in a random and equal way. See "Overview of the "Implicit Load Distribution" method" on page 186 and "Configuring the "Implicit Load Distribution" method" on page 191.

Implicit MEP is supported, if the Security Gateways with overlapping encryption domains are in the same community. If they are located in different communities, only one of the Security Gateways will be used for this encryption domain.

Overview of the "Implicit First to Respond" method

When there is no primary Security Gateway, all Security Gateways share "equal priority".

When all Security Gateways share "equal priority":

- Remote VPN peers send RDP/DPD packets to all the Security Gateways in the MEP configuration.
- The first Security Gateway to respond to the probing RDP/DPD packets gets chosen as the entry point to network.

The idea behind *first to respond* is "proximity". The Security Gateway which is "closer" to the remote VPN peer responds first.


- A VPN tunnel is opened with the first to respond. All subsequent connections pass through the chosen Security Gateway.
- If the Security Gateway ceases to respond, a new Security Gateway is chosen.

In a star VPN community, RDP/DPD packets are sent to the Security Gateways and the first to respond is used for routing only when:

1. There is more than one center Security Gateway
2. One of the following VPN routing options was selected:
 - **To center and to other satellites through center**
 - **To center, or through the center to other satellites, to internet and other VPN targets**

This setting is found on the **Community Properties > VPN Advanced > VPN Routing** page.

In this example scenario:



- MEP is **not** enabled on the VPN community
- First to respond method is used
- Security Gateway X accesses VPN domain A through Security Gateway A
- Security Gateway X accesses VPN domain B through Security Gateway B
- Security Gateway X accesses VPN domain C through Security Gateway A or B

Overview of the "Implicit Primary-Backup Security Gateways" method

Backup Security Gateways provide redundancy for primary Security Gateways.

The first Security Gateway is configured as the "primary," and the second Security Gateway as the "backup." If the primary Security Gateway fails, for whatever reason, the remote VPN peer detects that the link has gone down and works through the backup Security Gateway. The backup Security Gateway inherits the complete VPN domain of the primary Security Gateway. Failover within an existing connection is not supported; the current connection is lost.

When the primary Security Gateway is restored, new connections go through the primary Security Gateway, while connections that already exist will continue to work through the backup Security Gateway.

Important - When you use the Primary-Backup Security Gateways method, the encryption domains should not overlap.

Overview of the "Implicit Load Distribution" method

To prevent any one Security Gateway from being flooded with connections, the connections can be evenly shared amongst all the Security Gateways to distribute the load. When all Security Gateways share equal priority (no primary) and are MEP to the *same* VPN domain, it is possible to enable load distribution between the Security Gateways. The Security Gateways are probed with RDP/DPD packets, as in all other MEP configurations, to create a list of responding Security Gateways. A Security Gateway is randomly chosen from the list of responding Security Gateways. If a Security Gateway stops responding, a new Security Gateway is (randomly) chosen.

A new Security Gateway is randomly selected for every source/destination IP addresses pair. While the source and destination IP addresses remain the same, the connection continues through the chosen Security Gateway.

Routing Return Packets

To make sure return packets are routed correctly, the MEP Security Gateway can make use of either of these:

- IP Pool NAT (Static NAT)
- Route Injection Mechanism (RIM)

IP Pool NAT

IP Pool NAT is a type of NAT, in which source IP addresses from remote VPN domains are mapped to an IP address drawing from a pool of registered IP addresses. In order to maintain symmetric sessions with MEP Security Gateways, the MEP Security Gateway performs NAT with a range of IP addresses dedicated to that specific Security Gateway and should be routed within the internal network to the originating Security Gateway. When the returning packets reach the Security Gateway, the Security Gateway restores the original source IP address and forwards the packets to the source.

Route Injection Mechanism

Route Injection Mechanism (RIM) enables a Security Gateway to use a dynamic routing protocol to propagate the encryption domain of a VPN peer Security Gateway to the internal network. When a VPN tunnel is created, RIM updates the local routing table of the Security Gateway to include the encryption domain of the VPN peer.

When a tunnel to a MEP Security Gateway goes down, the Security Gateway removes the applicable "return route" from its own local routing table. This change is then distributed backwards to the routers behind the Security Gateway.

RIM is based both on the ability of the Security Gateway to update its local routing table, and the presence of a dynamic routing protocol to distribute the change to the network behind the Security Gateway. There is little sense in enabling RIM on the Security Gateway if a dynamic routing protocol is not available to distribute changes.

When MEP is enabled, RIM can be enabled only if permanent tunnels are enabled for the whole community. In a MEP configuration RIM is available when you use the **First to Respond**, **Manual set priority list**, and **VPN Domain** mechanisms. In the first two options, satellite Security Gateways "see" the center Security Gateways as unified as if one tunnel is connecting them. As a result, only the chosen MEP Security Gateway will inject the routes. In **VPN Domain** MEP, it could be that all MEP Security Gateways will inject the routes, which requires configuring the routers behind the MEP Security Gateways to return packets to the correct Security Gateway.

RIM is not available when **Random Selection** is the selected entry point mechanism.

For more information, see ["Route Injection Mechanism" on page 154](#).

Special Considerations

1. If one of the central Security Gateways is an externally managed Security Gateway:
 - The VPN domain of the central Security Gateways will not be automatically inherited by an externally managed Security Gateway
 - The RIM configuration will not be automatically downloaded
2. DAIP Security Gateways require DNS resolving in order to be configured as MEP Security Gateways.

Configuring MEP

To configure MEP, decide on:

1. The MEP method:
 - Explicit MEP.
 - Implicit MEP.
2. If required, method for returning reply packets:
 - IP Pool NAT
 - Route Injection Mechanism (see ["Route Injection Mechanism" on page 154](#)).

Configuring Explicit MEP

Explicit MEP is only available in Site-to-Site Star VPN communities where multiple center Security Gateways are defined.

To configure MEP:

1. In SmartConsole, click **Objects** menu > **Object Explorer**.
2. From the left tree, select **VPN Communities**.
3. Open the Star VPN Community object.
4. From the left tree, click **MEP**.
5. Select **Enable center gateways as MEP**.
6. Select the applicable entry point mechanism:
 - **First to respond**
 - **By VPN domain**
 - **Random selection**
 - **Manual priority list**

Notes:

- If you select **By VPN domain** or **Manually set priority list**, then in the **Advanced** section choose **First to respond** or **Random selection** to resolve how more than one Security Gateway with equal priority should be selected.
 - If you select **Manually set priority list**, then click **Set** to create a series of MEP rules.
7. Select a **Tracking** option, if required.
 8. Click **OK**.
 9. Install the Access Control Policy.

Configuring Implicit MEP

Configuring the "Implicit Primary-Backup" method

Configure the VPN Domain that includes the Primary Security Gateway and another VPN Domain that includes only the Backup Security Gateways.

Configure each Security Gateway as either the Primary Security Gateway, or a Backup Security Gateway.

Procedure:

Step	Instructions
1	Enable the Backup Gateway options in Global properties: 1. Click Menu > Global properties > VPN > Advanced. 2. Select Enable Backup Gateway. 3. Click OK to close the Global properties window.
2	Configure a Network Group object: 1. Click Objects menu > Object Explorer. 2. At the top, click New > Network Group. 3. Configure this Network Group object to contain only Backup Security Gateways. 4. Click OK to close the Network Group object window.
3	Configure the Primary Security Gateway object: 1. Click Objects menu > Object Explorer. 2. From the left tree, select Network Objects > Gateways & Servers. 3. Open the Primary Security Gateway object. 4. Click IPsec VPN. 5. Select Use Backup Gateways. 6. In the drop-down menu, select the Network Group object that contains the Backup Security Gateways. 7. Click OK to close the Primary Security Gateway object. This Security Gateway is now the Primary Security Gateway for this VPN domain.
4	Define the VPN for the Backup Security Gateways. Backup Security Gateways do not always have a VPN Domain of their own. They simply back up the Primary Security Gateway.

Step	Instructions
5	If the Backup Security Gateway does not have a VPN Domain of its own, the VPN Domain should include only the Backup Security Gateway itself: 1. Click Objects menu > Object Explorer. 2. From the left tree, select Network Objects > Gateways & Servers. 3. Open the Backup Security Gateway object. 4. Click Network Management > VPN Domain. 5. Select Manually defined. 6. Click the [...] button. 7. Select the Network Group object that contains only the Backup Security Gateways. 8. Click OK to close the Backup Security Gateway object. 9. Install the Access Control Policy on the Backup Security Gateways.
6	If the Backup Security Gateway does have a VPN Domain of its own: 1. Click Objects menu > Object Explorer. 2. From the left tree, select Network Objects > Gateways & Servers. 3. Open the Backup Security Gateway object. 4. Click Network Management > VPN Domain. 5. Make sure that the IP address of the Backup Security Gateway is not included in the VPN Domain of the Primary Security Gateway. 6. For each Backup Security Gateway, define a VPN Domain that does not overlap with the VPN Domain of other Backup Security Gateways. 7. Click OK to close the Backup Security Gateway object. 8. Install the Access Control Policy on the Backup Security Gateways. Important - There must be no overlap between the VPN Domain of the Primary Security Gateway and the VPN Domain of the Backup Security Gateways - no IP address can belong to both VPN Domains.

Configuring the "Implicit Load Distribution" method

To configure implicit MEP for random Security Gateway selection in SmartConsole:

Step	Instructions
1	Enable load distribution for MEP configurations: 1. Click Menu > Global properties. 2. From the left tree, click VPN > Advanced. 3. Select Enable load distribution for Multiple Entry Point configurations (Site to Site connections). 4. Click OK to close the Global properties window.
2	Configure a Network Group object: 1. Click Objects menu > Object Explorer. 2. At the top, click New > Network Group. 3. Configure this Network Group object to contain all the Security Gateways. 4. Click OK to close the Network Group object window.
3	Define the same VPN Domain for all the Security Gateways: 1. Click Objects menu > Object Explorer. 2. From the left tree, select Network Objects > Gateways & Servers. 3. Open each Security Gateway object. 4. Click Network Management > VPN Domain. 5. Select Manually defined. 6. Click the [...] button. 7. Select the Network Group object that contains all the Security Gateways. 8. Click OK to close the Security Gateway object.
4	Install the Access Control Policy on all the Security Gateways.

Configuring IP Pool NAT

To configure IP Pool NAT for Site to Site VPN in SmartConsole:

Step	Instructions
1	Enable IP Pool NAT in the Global Properties: 1. Click Menu > Global properties. 2. Click NAT - Network Address Translation. 3. Select Enable IP Pool NAT. 4. Select the applicable options (None, Log, or Alert) for: ▪ Address exhaustion track ▪ Address allocation and release track 5. Click OK to close the Global properties window.
2	For <i>each</i> Security Gateway, configure an object that represents the IP Pool NAT addresses for that Security Gateway: 1. Click Objects menu > Object Explorer. 2. The object that represents the IP Pool NAT addresses can be one of these objects: ▪ Network - At the top, click New > Network ▪ Network Group - At the top, click New > Network Group ▪ Address Range - At the top, click New > Network Object > Address Range > Address Range. 3. Configure this object to contain the applicable IP addresses. 4. Click OK to close the object with IP Pool NAT addresses.

Step	Instructions
3	In <i>each</i> Security Gateway, configure IP Pool NAT settings: 1. Click Objects menu > Object Explorer. 2. From the left tree, select Network Objects > Gateways & Servers. 3. Open <i>each</i> Security Gateway object. 4. Click NAT > IP Pool NAT. 5. Select one of these two options: ▪ Allocate IP Addresses from. If you choose this option, then select the object that represents the IP Pool NAT addresses for that Security Gateway. ▪ Define IP Pool addresses on Gateway interfaces. If you choose this option, then you must configure the IP Pool NAT on <i>each</i> required interface: (i) From the left tree, click Network Management. (ii) Edit each required interface. (iii) From the left tree, click General. (iv) In the Topology section, click Modify. (v) In the IP Pool NAT section, select the object that represents the IP Pool NAT addresses. (vi) Click OK. 1. Select the applicable options: ▪ Use IP Pool NAT for VPN clients connections ▪ Use IP Pool NAT for gateway to gateway connections ▪ Prefer IP Pool NAT over Hide NAT 2. Click Advanced to configure the advanced IP Pool NAT settings. Click OK. 3. Click OK to close the Security Gateway object.
4	Install the Access Control Policy on all Security Gateways.
5	Edit the routing table for each internal router, so that packets with an IP address assigned from the IP Pool NAT are routed to the applicable Security Gateway.

Resolving Connectivity Issues

IPsec NAT-Traversal

NAT-T (NAT traversal or UDP encapsulation) makes sure that IPsec VPN connections stay open when traffic goes through Security Gateways or devices that use NAT.

When an IP packet passes through a network address translator device, it is changed in a way that is not compatible with IPsec. To protect the original IPsec encoded packet, NAT traversal encapsulates it with an additional layer of UDP and IP headers.

For IPsec to work with NAT traversal, these protocols must be allowed through the NAT interface(s):

- IKE - UDP port 500
- IPsec NAT-T - UDP port 4500
- Encapsulating Security Payload (ESP) - IP protocol number 50
- Authentication Header (AH) - IP protocol number 51

Configuring NAT-Traversal

To configure NAT-T for Site to Site VPN:

1. In SmartConsole, from the left navigation panel, click **Gateways & Servers**.
2. Open the applicable Security Gateway object with enabled IPsec VPN Software Blade.
3. From the left tree, click **IPsec VPN > VPN Advanced**.
4. Make sure to select **Support NAT traversal (applies to Remote Access and Site to Site connections)**.

NAT-Traversal is enabled by default when a NAT device is detected.

5. Click **OK**.
6. Install the Access Control Policy.

Advanced NAT-T Configuration

These kernel parameters are defined for each Security Gateway and control NAT-T for Site to Site VPN:

Item	Description	Default Value
<code>offer_nat_t_initiator</code>	Initiator sends NAT-T traffic. For more information, see sk177823 .	true
<code>offer_nat_t_responder_for_known_gw</code>	Responder accepts NAT-T traffic from known Security Gateways.	true
<code>force_nat_t</code>	Force NAT-T, even if there is no NAT-T device. For more information, see sk166037 . Important - The value of this parameter must be the same for all VPN peers.	false

You can edit these kernel parameters in the Database Tool (GuiDBEdit Tool) (see [sk13009](#)):

1. (Recommended) back up the Security Management Server / Domain Management Server.
2. Close all SmartConsole windows.
3. Connect with Database Tool (GuiDBEdit Tool) to the Security Management Server / Domain Management Server.
4. In the upper left pane, go to **Table - Network Objects - network_objects**.
5. In the upper right pane, select the relevant Security Gateway / Cluster object.
6. Press **CTRL + F** (or go to **Search** menu - **Find**) - paste - NAME OF PARAMETER click **Find Next**.
7. In the lower pane, right click NAME OF PARAMETER- select **Edit...** - set the desired value.

Command Line Reference

VPN commands generate status information regarding VPN processes, or are used to stop and start specific VPN services.

All VPN commands are executed on the Security Gateway and Cluster Members.

For more information, see the:

- [*R81.10 CLI Reference Guide*](#).
- [*R81.10 Remote Access VPN Administration Guide*](#).

This section contains a shortened list of applicable commands.

Syntax Legend

Whenever possible, this guide lists commands, parameters and options in the alphabetical order.

This guide uses this convention in the Command Line Interface (CLI) syntax:

Character	Description
TAB	Shows the available nested subcommands: <pre>main command → nested subcommand 1 → → nested subsubcommand 1-1 → → nested subsubcommand 1-2 → nested subcommand 2</pre> Example: <pre>cpwd_admin config -a <options> -d <options> -p -r del <options></pre> Meaning, you can run only one of these commands: ■ This command: <pre>cpwd_admin config -a <options></pre> ■ Or this command: <pre>cpwd_admin config -d <options></pre> ■ Or this command: <pre>cpwd_admin config -p</pre> ■ Or this command: <pre>cpwd_admin config -r</pre> ■ Or this command: <pre>cpwd_admin del <options></pre>
Curly brackets or braces { }	Enclose a list of available commands or parameters, separated by the vertical bar . User can enter only one of the available commands or parameters.

Character	Description
Angle brackets < >	Enclose a variable. User must explicitly specify a supported value.
Square brackets or brackets []	Enclose an optional command or parameter, which user can also enter.

ike debug

Background

Starting in R81.10, separate daemons handle different VPN connections:

- The VPN daemon `vpnd`.

Handles these VPN connections:

- Site-to-Site connections from peer Security Gateways with a Statically Assigned IP address
- All connections from non-IPsec Remote Access clients (SSL Network Extender)
- Multi-Portal traffic

Listens on these ports on a Security Gateway:

- IKE: 500 (UDP)
- IKE NAT-T: 4500 (UDP)
- TCPT: 444 (TCP)
- Session infrastructure manager: 9996 (TCP)

This process is a child of the FWD process (see the `$FWDIR/conf/fwauthd.conf` file on a Security Gateway).

- The IKE daemon `iked` (introduced in R81.10).

Handles these VPN connections:

- All connections from IKE Remote Access clients (for example, Endpoint clients)
- Site-to-Site connections from peer Security Gateways with a Dynamically Assigned IP address (DAIP)
- Large Scale VPN (LSV) connections
- Connections from SmartLSM ROBO gateways

Listens on these ports on a Security Gateway:

- IKE: 30500 (UDP)
- IKE NAT-T: 34500 (UDP)
- Session infrastructure manager: 9994 (TCP)
- L2TP: 1701 (UDP)

Syntax: `vpn iked`

This process is a child of the FWD process (see the `$FWDIR/conf/fwauthd.conf` file on a Security Gateway).

- The CCC daemon `cccd` (introduced in R81.10).

Responsible for the Circuit Cross-Connect (CCC) protocol, while:

- IKE for the same clients runs in the IKE daemon `iked`
- CCC TLS for the same clients runs in the VPN daemon `vpnd`

Listens on these ports on a Security Gateway:

- TCPT: 444 (TCP)
- Session infrastructure manager: 9993 (TCP)

Syntax: `vpn cccd`

This process is a child of the FWD process (see the `$FWDIR/conf/fwauthd.conf` file on a Security Gateway).

Disabling the IKE daemon "iked"

It is possible to configure the Security Gateway to disable the IKE daemon `iked` and work in the legacy mode, in which the VPN daemon `vpnd` handles all VPN connections.

Procedure

Step	Instructions
1	Connect to the command line on the Security Gateway / each Cluster Member.
2	Log in to the Expert mode.
3	Examine the current status: <pre>vpniked status</pre> Possible outputs: ■ <code>vpn: 'iked' is enabled</code> Means the IKE daemon <code>iked</code> is enabled (new mode in R81.10 and higher). This is the default. Proceed to the next step. ■ <code>vpn: 'iked' is disabled</code> Means the IKE daemon <code>iked</code> is disabled (legacy mode as in R81 and lower). There is nothing to do.
4	Disable the IKE daemon: <pre>vpniked disable</pre> Important ■ This command makes the required changes in the <code>\$FWDIR/boot/modules/fwkernel.conf</code> file. ■ This command installs the local policy, which can disconnect your SSH session. ■ This change survives reboot.

To enable the IKE daemon `iked` after you disable it, run this command:

```
vpniked enable
```

Important

- This command makes the required changes in the `$FWDIR/boot/modules/fwkernel.conf` file.
- This command installs the local policy, which can disconnect your SSH session.
- This change survives reboot.

Disabling the CCC daemon "cccd"

It is possible to configure the Security Gateway to disable the CCC daemon `cccd` and work in the legacy mode, in which the VPN daemon `vpnd` handles all VPN connections.

Procedure

Step	Instructions
1	Connect to the command line on the Security Gateway / each Cluster Member.
2	Log in to the Expert mode.
3	Examine the current status: <pre>vpn cccd status</pre> Possible outputs: ▪ <code>vpn: 'cccd' is enabled</code> Means the CCC daemon <code>cccd</code> is enabled (new mode in R81.10 and higher). This is the default. Proceed to the next step. ▪ <code>vpn: 'cccd' is disabled</code> Means the CCC daemon <code>cccd</code> is disabled (legacy mode as in R81 and lower). There is nothing to do.
4	Disable the CCC daemon: <pre>vpn cccd disable</pre>  Important ▪ This command makes the required changes in the <code>\$FWDIR/boot/modules/fwkernel.conf</code> file. ▪ This command installs the local policy, which can disconnect your SSH session. ▪ This change survives reboot.

To enable the CCC daemon `cccd` after you disable it, run this command:

```
vpn cccd enable
```

 Important

- This command makes the required changes in the `$FWDIR/boot/modules/fwkernel.conf` file.
- This command installs the local policy, which can disconnect your SSH session.
- This change survives reboot.

Description

This "ike debug" command instructs the IKE daemon `iked` to write debug messages to these log files:

File Description	IKED Log File
Main debug output file	<code>\$FWDIR/log/iked.elg</code>
IKEv1 output	<code>\$FWDIR/log/ike.elg</code>
IKEv2 output	<code>\$FWDIR/log/ikev2.xml</code>

Debugging of the IKE daemon `iked` is based on Debug Topics and Debug Levels:

- A Debug Topic is a specific area, on which to perform debugging.
Check Point Support provides the specific Debug Topics when needed.
- Debug Levels range from 1 (least informative) to 5 (most informative - write all debug messages).

In addition, see ["vpn debug" on page 217](#).

Syntax

```
ike debug
  on [<Debug_Topic>=<Debug_Level>]
  off
  ikeon [-s <Size_in_MB>]
  ikeoff
  trunc [<Debug_Topic>=<Debug_Level>]
  truncon [<Debug_Topic>=<Debug_Level>]
  truncoff
  timeon [<Seconds>]
  timeoff
  ikefail [-s <Size_in_MB>]
  mon
  moff
  say ["String"]
  tunnel [<Level>]
```

Parameters

Parameter	Description
No Parameters	Shows the built-in usage.
on	Turns on high level IKE debug. The debug writes the information in these files: ▪ \$FWDIR/log/ike.elg* ▪ \$FWDIR/log/iked.elg*
<Debug_Topic>=<Debug_Level>	Specifies the Debug Topic and the Debug Level. Check Point Support provides these. ★ Best Practice - Run this command to start the debug: <pre>ike debug trunc ALL=5</pre>
off	Turns off all IKE debug. ★ Best Practice - Run one of these commands to stop the IKE debug: <pre>ike debug off</pre> <pre>ike debug truncoff</pre>

Parameter	Description
ikeon [-s <Size_in_MB>]	Turns on the IKE trace. The debug writes the information in the <code>\$FWDIR/log/ike.elg*</code> files. You can specify the size of the <code>\$FWDIR/log/ike.elg</code> file, when to perform the log rotation (close the current active file, rename it, open a new active file).
ikeoff	Turns off the IKE trace. Run this command to stop the IKE trace: <pre>ike debug ikeoff</pre>
trunc or truncon	This command: 1. Rotates the <code>\$FWDIR/log/ike.elg</code> file 2. Rotates the <code>\$FWDIR/log/iked.elg</code> file 3. Rotates the <code>\$FWDIR/log/ikev2.xml1</code> file 4. Starts the IKE daemon debug  Best Practice - Run this command to start the IKE debug: <pre>ike debug trunc ALL=5</pre>
truncoff	Stops the IKE daemon debug. Run one of these commands to stop the IKE debug: <pre>ike debug truncoff</pre> <pre>ike debug off</pre>
timeon [<Seconds>]	Enables the timestamp in the log files. Prints one timestamp after the specified number of seconds. By default, prints the timestamp every 10 seconds.
timeoff	Disables the timestamp in the log files every number of seconds.
ikefail [-s <Size_in_MB>]	Logs failed IKE negotiations. You can specify the size of the <code>\$FWDIR/log/ike.elg</code> file, when to perform the log rotation (close the current active file, rename it, open a new active file).
mon	Enables the IKE Monitor. Writes the IKE packets in the <code>\$FWDIR/log/ikemonitor.snoop</code> file.  Warning - The output file may contain user X-Auth passwords. Make sure the file is protected.
moff	Disables the IKE Monitor.

Parameter	Description
say "String"	Writes the specified text string in the <code>\$FWDIR/log/iked.elg</code> file. For example, run: <code>ike debug say "BEGIN TEST"</code>  Notes: - Run this command after you start the VPN debug (with one of these commands: <code>"ike debug on"</code>, <code>"ike debug trunc"</code>, or <code>"ike debug truncon"</code>). - The length of the string is limited to 255 characters.
tunnel [<Debug_Level>]	This command: - Rotates the <code>\$FWDIR/log/ike.elg</code> file - Starts the IKE daemon debug with these two Debug Topics: tunnel ikev2 If the <Debug_Level> is 2,3,4 or 5, then also enables this Debug Topic: CRLCache The debug writes the information in these files: <code>\$FWDIR/log/ike.elg*</code> <code>\$FWDIR/log/iked.elg*</code>

Return Values

- 0 (zero) for success
- any other value for failure (typically, -1 or 1)

vpn

Description

Configures VPN settings.

Shows VPN information.

Syntax

```
vpn
    check_ttm
    compreset
    compstat
    crl_zap
    crlview
    debug
    dll
    drv
    dump_psk
    ipafile_check
    ipafile_users_capacity
    macutil
    mep_refresh
    neo_proto
    nssm_topology
    overlap_encdom
    rim_cleanup
    rll
    set_slim_server
    set_snx_encdom_groups
    set_trac
    shell
    show_tcpt
    sw_topology
    {tunnelutil | tu}
    ver
```

Parameters

Parameter	Description
check_ttm	Makes sure the specified TTM file is valid. See "vpn check_ttm" on page 211 .
compreset	Resets compression and decompression statistics counters. See "vpn compreset" on page 212 .
compstat	Shows compression and decompression statistics counters. See "vpn compstat" on page 213 .
crl_zap	Erases all Certificate Revocation Lists (CRLs) from the cache. See "vpn crl_zap" on page 214 .
crlview	Retrieves the Certificate Revocation List (CRL) from various distribution points and shows it for the user. See "vpn crlview" on page 215 .
debug	Controls the debug of <code>vpnd</code> daemon and IKE. See "vpn debug" on page 217 .
dll	Works with DNS Lookup Layer. See "vpn dll" on page 227 .
drv	Controls the VPN kernel module. See "vpn drv" on page 228 .
dump_psk	Shows hash (SHA256) of peers' pre-shared-keys. See "vpn dump_psk" on page 229 .
ipafile_check	Verifies a candidate for the <code>\$FWDIR/conf/ipassignment.conf</code> file. See "vpn ipafile_check" on page 230 .
ipafile_users_capacity	Shows and configures the capacity in the <code>\$FWDIR/conf/ipassignment.conf</code> file. See "vpn ipafile_users_capacity" on page 231 .
macutil	Shows a generated MAC address for each user name when you use Remote Access VPN with Office Mode. See "vpn macutil" on page 232 .
mep_refresh	Initiates MEP re-decision. See "vpn mep_refresh" on page 233 .

Parameter	Description
neo_proto	Controls the NEO client protocol. See "vpn neo_proto" on page 234 .
nssm_topology	Generates and uploads a topology in NSSM format to an NSSM server. See "vpn nssm_topology" on page 235 .
overlap_encdom	Shows all overlapping VPN domains. See "vpn overlap_encdom" on page 236 .
rim_cleanup	Cleans RIM routes. See "vpn rim_cleanup" on page 238 .
rll	Works with Route Lookup Layer. See "vpn rll" on page 239 .
set_slim_server	Deprecated. See "vpn set_slim_server" on page 240 .
set_snx_encdom_groups	Controls the encryption domain per usergroup feature for SSL Network Extender. See "vpn set_snx_encdom_groups" on page 241 .
set_trac	Controls the TRAC server. See "vpn set_trac" on page 242 .
shell	VPN Command Line Interface. See "vpn shell" on page 243 .
show_tcpt	Shows Visitor Mode users. See "vpn show_tcpt" on page 250 .
sw_topology	Downloads the topology for a UTM-1 Edge or Safe@Office device.  Note - R81.10 does not support UTM-1 Edge and Safe@Office devices. The information about this command is provided only to describe the existing syntax option until it is removed completely. See "vpn sw_topology" on page 251 .
tunnelutil tu	Launches the TunnelUtil tool, which is used to control VPN tunnels. See "vpn tu" on page 252 .
ver	Shows the major version number and build number of the VPN kernel module. See "vpn ver" on page 262 .

vpn check_ttm

Description

Makes sure the specified TTM file contains valid syntax.

Syntax

```
vpn check_ttm <Path to TTM file>
```

Parameters

Parameter	Description
<Path to TTM file>	Specifies the full path and name of the TTM file.

Example

```
[Expert@MyGW:0]# find / -name \*.ttm -type f
find: /proc/64899: No such file or directory
/var/opt/CPsuite-R81.10/fw1/conf/fw_client_1.ttm
/var/opt/CPsuite-R81.10/fw1/conf/nemo_client_1.ttm
/var/opt/CPsuite-R81.10/fw1/conf/neo_client_1.ttm
/var/opt/CPsuite-R81.10/fw1/conf/iphone_client_1.ttm
/var/opt/CPsuite-R81.10/fw1/conf/topology_trans_tmpl.ttm
/var/opt/CPsuite-R81.10/fw1/conf/vpn_client_1.ttm
/var/opt/CPsuite-R81.10/fw1/conf/trac_client_1.ttm
... ..
[Expert@MyGW:0]#
[Expert@MyGW:0]# vpn check_ttm /var/opt/CPsuite-R81.10/fw1/conf/trac_client_1.ttm

Summary for the file: trac_client_1.ttm
    result: the file passed the check without any problems

[Expert@MyGW:0]#
```

vpn compreset

Description

Resets compression and decompression statistics counters.

Syntax

```
vpn compreset
```

Example

```
[Expert@MyGW:0]# vpn compreset  
Compression statistics were reset.  
[Expert@MyGW:0]#
```

vpn compstat

Description

Shows compression and decompression statistics counters.

Syntax

```
vpn compstat
```

Example

```
[Expert@MyGW:0]# vpn compstat

Compression: sum of all instances :

Compression:
=====
Bytes before compression          : 0
Bytes after compression           : 0
Compression overhead (bytes)      : 0
Bytes that were not compressed    : 0
Compressed packets                : 0
Packets that were not compressed  : 0
Compression errors                : 0

Pure compression ratio            : 0.000000
Effective compression ratio       : 0.000000

Decompression:
=====
Bytes before decompression         : 0
Bytes after decompression          : 0
Decompression overhead (bytes)    : 0
Decompressed packets              : 0
Decompression errors              : 0
Pure decompression ratio           : 0.000000
[Expert@MyGW:0]#
```

vpn crl_zap

Description

Erases all Certificate Revocation Lists (CRLs) from the cache.

Syntax

```
vpn crl_zap
```

Return Values

- 0 (zero) for success
- any other value for failure

vpn crlview

Description

Retrieves the Certificate Revocation List (CRL) from various distribution points and shows it for the user.

Syntax

```
vpn crlview [-d]
             -obj <Network Object Name> -cert <Certificate Object Name>
             -f <Certificate File>
             -view
```

Parameters

Parameter	Description
-d	Runs the command in debug mode. Use only if you troubleshoot the command itself.  Best Practice - If you use this parameter, then redirect the output to a file, or use the script command to save the entire CLI session.
-obj <Network Object Name>	Specifies the name of the CA network object.
-cert <Certificate Object Name>	Specifies the name of the certificate object.
-f <Certificate File>	Specifies the path and the name of the certificate file.
-view	Shows the CRL.

Return Values

- 0 (zero) for success
- any other value for failure

Example 1

```
vpn crlview -obj <MyCA> -cert <MyCert>
```

1. The VPN daemon contacts the Certificate Authority called **MyCA** and locates the certificate called **MyCert**.
2. The VPN daemon extracts the certificate distribution point from the certificate.
3. The VPN daemon goes to the distribution point and retrieves the CRL. The distribution point can be an LDAP or HTTP server.
4. The VPN daemon shows it to the standard output.

Example 2

```
vpn crlview -f /var/log/MyCert
```

1. The VPN daemon extracts the certificate distribution point from the certificate file called **MyCert**.
2. The VPN daemon goes to the distribution point and retrieves the CRL. The distribution point can be an LDAP or HTTP server.
3. The VPN daemon shows the CRL to the standard output.

Example 3

```
vpn crlview -view <Lastest CRL>
```

If the CRL was retrieved in the past, this command instructs the VPN daemon to show the contents to the standard output.

vpn debug

Background

Starting in R81.10, separate daemons handle different VPN connections:

- The VPN daemon `vpnd`.

Handles these VPN connections:

- Site-to-Site connections from peer Security Gateways with a Statically Assigned IP address
- All connections from non-IPsec Remote Access clients (SSL Network Extender)
- Multi-Portal traffic

Listens on these ports on a Security Gateway:

- IKE: 500 (UDP)
- IKE NAT-T: 4500 (UDP)
- TCPT: 444 (TCP)
- Session infrastructure manager: 9996 (TCP)

This process is a child of the FWD process (see the `$FWDIR/conf/fwauthd.conf` file on a Security Gateway).

- The IKE daemon `iked` (introduced in R81.10).

Handles these VPN connections:

- All connections from IKE Remote Access clients (for example, Endpoint clients)
- Site-to-Site connections from peer Security Gateways with a Dynamically Assigned IP address (DAIP)
- Large Scale VPN (LSV) connections
- Connections from SmartLSM ROBO gateways

Listens on these ports on a Security Gateway:

- IKE: 30500 (UDP)
- IKE NAT-T: 34500 (UDP)
- Session infrastructure manager: 9994 (TCP)
- L2TP: 1701 (UDP)

Syntax: `vpn iked`

This process is a child of the FWD process (see the `$FWDIR/conf/fwauthd.conf` file on a Security Gateway).

- The CCC daemon `cccd` (introduced in R81.10).

Responsible for the Circuit Cross-Connect (CCC) protocol, while:

- IKE for the same clients runs in the IKE daemon `iked`
- CCC TLS for the same clients runs in the VPN daemon `vpnd`

Listens on these ports on a Security Gateway:

- TCPT: 444 (TCP)
- Session infrastructure manager: 9993 (TCP)

Syntax: `vpn cccd`

This process is a child of the FWD process (see the `$FWDIR/conf/fwauthd.conf` file on a Security Gateway).

Disabling the IKE daemon "iked"

It is possible to configure the Security Gateway to disable the IKE daemon `iked` and work in the legacy mode, in which the VPN daemon `vpnd` handles all VPN connections.

Procedure

Step	Instructions
1	Connect to the command line on the Security Gateway / each Cluster Member.
2	Log in to the Expert mode.
3	Examine the current status: <pre>vpniked status</pre> Possible outputs: ■ <code>vpn: 'iked' is enabled</code> Means the IKE daemon <code>iked</code> is enabled (new mode in R81.10 and higher). This is the default. Proceed to the next step. ■ <code>vpn: 'iked' is disabled</code> Means the IKE daemon <code>iked</code> is disabled (legacy mode as in R81 and lower). There is nothing to do.
4	Disable the IKE daemon: <pre>vpniked disable</pre> Important ■ This command makes the required changes in the <code>\$FWDIR/boot/modules/fwkernel.conf</code> file. ■ This command installs the local policy, which can disconnect your SSH session. ■ This change survives reboot.

To enable the IKE daemon `iked` after you disable it, run this command:

```
vpniked enable
```

Important

- This command makes the required changes in the `$FWDIR/boot/modules/fwkernel.conf` file.
- This command installs the local policy, which can disconnect your SSH session.
- This change survives reboot.

Disabling the CCC daemon "cccd"

It is possible to configure the Security Gateway to disable the CCC daemon `cccd` and work in the legacy mode, in which the VPN daemon `vpnd` handles all VPN connections.

Procedure

Step	Instructions
1	Connect to the command line on the Security Gateway / each Cluster Member.
2	Log in to the Expert mode.
3	Examine the current status: <pre>vpn cccd status</pre> Possible outputs: ▪ <code>vpn: 'cccd' is enabled</code> Means the CCC daemon <code>cccd</code> is enabled (new mode in R81.10 and higher). This is the default. Proceed to the next step. ▪ <code>vpn: 'cccd' is disabled</code> Means the CCC daemon <code>cccd</code> is disabled (legacy mode as in R81 and lower). There is nothing to do.
4	Disable the CCC daemon: <pre>vpn cccd disable</pre>  Important ▪ This command makes the required changes in the <code>\$FWDIR/boot/modules/fwkernel.conf</code> file. ▪ This command installs the local policy, which can disconnect your SSH session. ▪ This change survives reboot.

To enable the CCC daemon `cccd` after you disable it, run this command:

```
vpn cccd enable
```

Important

- This command makes the required changes in the `$FWDIR/boot/modules/fwkernel.conf` file.
- This command installs the local policy, which can disconnect your SSH session.
- This change survives reboot.

Description

This "vpn debug" command instructs the VPN daemon `vpnd` to write debug messages to these log files:

File Description	VPND Log File	IKED Log File	CCCD Log File
Main debug output file	<code>\$FWDIR/log/vpnd.elg*</code>	<code>\$FWDIR/log/iked.elg</code>	<code>\$FWDIR/log/cccd.elg</code>
IKEv1 output	<code>\$FWDIR/log/legacy_ike.elg</code>	<code>\$FWDIR/log/ike.elg</code>	N / A
IKEv2 output	<code>\$FWDIR/log/legacy_ikev2.xml</code>	<code>\$FWDIR/log/ikev2.xml</code>	N / A

Debugging of the VPN daemon `vpnd` is based on Debug Topics and Debug Levels:

- A Debug Topic is a specific area, on which to perform debugging.
For example, if the Debug Topic is `LDAP`, all traffic between the VPN daemon and the LDAP server is written to the log file.
Check Point Support provides the specific Debug Topics when needed.
- Debug Levels range from 1 (least informative) to 5 (most informative - write all debug messages).

For more information, see [sk180488](#).

In addition, see *"ike debug" on page 199*.

Syntax

```
vpn debug
  on [<Debug_Topic>=<Debug_Level>]
  off
  ikeon [-s <Size_in_MB>]
  ikeoff
  trunc [<Debug_Topic>=<Debug_Level>]
  truncon [<Debug_Topic>=<Debug_Level>]
  truncoff
  timeon [<Seconds>]
  timeoff
  ikfail [-s <Size_in_MB>]
  mon
  moff
  say ["String"]
  tunnel [<Level>]
```

Parameters

Parameter	Description
No Parameters	Shows the built-in usage.
on	Turns on high level VPN debug. The debug writes the information in the <code>\$FWDIR/log/vpnd.elg*</code> files. Note - If you disabled the IKE daemon <code>iked</code> to work in the legacy mode, then the debug also writes the information in the <code>\$FWDIR/log/legacy_ike.elg*</code> files.
<Debug_Topic>=<Debug_Level>	Specifies the Debug Topic and the Debug Level. Check Point Support provides these. ★ Best Practice - Run this command to start the VPN debug: <pre>vpn debug trunc ALL=5</pre>
off	Turns off the VPN debug (in legacy mode, also turns off the IKE debug). ★ Best Practice - Run one of these commands to stop the VPN daemon debug: <pre>vpn debug off</pre> <pre>vpn debug truncoff</pre>

Parameter	Description
ikeon [-s <Size_in_MB>]	Note - Applies only if you disabled the IKE daemon <code>iked</code> to work in the legacy mode. Turns on the IKE trace. The debug writes the information in the <code>\$FWDIR/log/legacy_ike.elg*</code> files. You can specify the size of the <code>\$FWDIR/log/legacy_ike.elg</code> file, when to perform the log rotation (close the current active file, rename it, open a new active file).
ikeoff	Note - Applies only if you disabled the IKE daemon <code>iked</code> to work in the legacy mode. Turns off the IKE trace. Run this command to stop the IKE trace: <pre>vpn debug ikeoff</pre>
trunc or truncon	This command: 1. Rotates the <code>\$FWDIR/log/vpnd.elg</code> file 2. If you disabled the IKE daemon <code>iked</code> to work in the legacy mode: a. Truncates the <code>\$FWDIR/log/legacy_ike.elg</code> file b. Truncates the <code>\$FWDIR/log/legacy_ikev2.xml11</code> file 3. Starts the VPN daemon <code>vpnd debug</code> 4. If you disabled the IKE daemon <code>iked</code> to work in the legacy mode: a. Starts the IKE debug  Best Practice - Run this command to start the VPN debug: <pre>vpn debug trunc ALL=5</pre>
truncoff	Stops the VPN daemon debug. Run one of these commands to stop the VPN debug: <pre>vpn debug truncoff</pre> <pre>vpn debug off</pre>
timeon [<Seconds>]	Enables the timestamp in the log files. Prints one timestamp after the specified number of seconds. By default, prints the timestamp every 10 seconds.
timeoff	Disables the timestamp in the log files every number of seconds.

Parameter	Description
ikefail [-s <Size_in_MB>]	Note - Applies only if you disabled the IKE daemon <code>iked</code> to work in the legacy mode. Logs failed IKE negotiations. You can specify the size of the <code>\$FWDIR/log/legacy_ike.elg</code> file, when to perform the log rotation (close the current active file, rename it, open a new active file).
mon	Note - Applies only if you disabled the IKE daemon <code>iked</code> to work in the legacy mode. Enables the IKE Monitor. Writes the IKE packets in the <code>\$FWDIR/log/ikemonitor.snoop</code> file.  Warning - The output file may contain user X-Auth passwords. Make sure the file is protected.
moff	Note - Applies only if you disabled the IKE daemon <code>iked</code> to work in the legacy mode. Disables the IKE Monitor.
say "String"	Writes the specified text string in the <code>\$FWDIR/log/vpnd.elg</code> file. For example, run: <code>vpn debug say "BEGIN TEST"</code>  Notes: - Run this command after you start the VPN debug (with one of these commands: <code>"vpn debug on"</code>, <code>"vpn debug trunc"</code>, or <code>"vpn debug truncon"</code>). - The length of the string is limited to 255 characters.
tunnel [<Debug_Level>]	This command: - Rotates the <code>\$FWDIR/log/vpnd.elg</code> file - Truncates the <code>\$FWDIR/log/legacy_ike.elg</code> file (if you disabled the IKE daemon <code>iked</code> to work in the legacy mode) - Starts the VPN daemon debug with these two Debug Topics: <pre>tunnel ikev2</pre> If the <code><Debug_Level></code> is 2,3,4 or 5, then also enables this Debug Topic: <pre>CRLCache</pre> - Starts the IKE debug (if you disabled the IKE daemon <code>iked</code> to work in the legacy mode)

Return Values

- 0 (zero) for success
- any other value for failure (typically, -1 or 1)

vpn dll

Description

Works with VPN DNS Lookup Layer:

- Save the DNS Lookup Layer information to the specified file.
- Resolve the specified hostname.

Syntax

```
vpn dll
    dump <File>
    resolve <HostName>
```

Parameters

Parameter	Description
dump <File>	Saves the DNS Lookup Layer information (DNS Names and IP Addresses) to the specified file.
resolve <HostName>	Resolves the specified hostname. The command saves the last specified hostname in this file: \$FWDIR/tmp/vpnd_cmd.tmp

vpn drv

Description

Controls the VPN kernel module.

Syntax

```
vpn drv {off | on | stat}
```

Parameters

Parameter	Description
off	Stops the VPN kernel module
on	Starts the VPN kernel module
stat	Shows the current status of the VPN kernel module

Example

```
[Expert@MyGW:0]# vpn drv stat
VPN-1 module active
[Expert@MyGW:0]#
```

vpn dump_psk

Description

Shows hash (SHA256) of peers' pre-shared-keys.

Syntax

```
vpn dump_psk
```

vpn ipafilename_check

Description

Verifies a candidate for the `$FWDIR/conf/ipassignment.conf` file.

Syntax

```
vpn ipafilename_check <File> [{err | warn | detail}] [verify_group_names]
```

Parameters

Parameter	Description
<code><File></code>	Specifies the full path and name of the candidate file.
<code>{err warn detail}</code>	Specifies the how much information to show about the candidate file: ▪ <code>err</code> - Only errors ▪ <code>warn</code> - Only warnings ▪ <code>detail</code> - All details
<code>verify_group_names</code>	Examines the group names.

vpn ipafire_users_capacity

Description

- Shows the current capacity in the `$FWDIR/conf/ipassignment.conf` file.
- Configures the new capacity in the `$FWDIR/conf/ipassignment.conf` file.

Syntax

```
vpn ipafire_users_capacity get
```

```
vpn ipafire_users_capacity set <128-32768>
```

Parameters

Parameter	Description
get	Shows the current capacity.
set <128-32768>	Configures the new capacity to the specified number of users.  Notes: ▪ The default is 1024 entries. ▪ This command configures the amount of memory reserved to store usernames.

Example

```
[Expert@MyGW:0]# vpn ipafire_users_capacity get
The gateway can currently read 1024 users from the ipassignment.conf file
[Expert@MyGW:0]#
```

vpn macutil

Description

Shows a generated MAC address for each user name when you use Remote Access VPN with Office Mode.

This command is applicable only when allocating IP addresses through DHCP.

Remote Access VPN users in Office Mode receive an IP address, which is mapped to a hardware or MAC address.

Syntax

```
vpn macutil <username>
```

Example

```
# vpn macutil John  
20-0C-EB-26-80-7D, "John"
```

vpn mep_refresh

Description

Initiates MEP re-decision.

Used in "backup stickiness" configuration to initiate MEP re-decision (fail back to primary Security Gateway, if possible).

Syntax

```
vpn mep_refresh
```

vpn neo_proto

Description

Controls the NEO client protocol.



Important - This command is for Check Point use only.

Syntax

```
vpn neo_proto {off | on}
```

Parameters

Parameter	Description
off	Disables the NEO client protocol.
on	Enables the NEO client protocol.

vpn nssm_topology

Description

Generates and uploads a topology in NSSM format to an NSSM server.

Syntax

```
vpn nssm_topology -url <"url"> -dn <"dn"> -name <"name"> -pass
<"password"> [-action {bypass | drop}] [-print_xml]
```

Parameters

Parameter	Description
-url <"url">	URL of the NSSM server.
-dn <"dn">	Distinguished Name of the NSSM server (needed to establish an SSL connection).
-name <"name">	Valid login name for the NSSM server.
-pass <"password">	Valid password for the NSSM server.
-action {bypass drop}	Specifies the action that the Symbian client should take, if the packet is not destined for an IP address in the VPN domain. Bypass is the default.
-print_xml	Writes the topology to a file in XML format.

vpn overlap_encdom

Description

Shows all overlapping VPN domains.

Some IP addresses might belong to two or more VPN domains.

The command alerts for overlapping encryption domains if one or both of the following conditions exist:

- The same VPN domain is defined for both Security Gateways.
- If the Security Gateway has multiple interfaces, and one or more of the interfaces has the same IP address and netmask.

Syntax

```
vpn overlap_encdom [communities | traditional]
```

Parameters

Parameter	Description
<code>communities</code>	Shows all pairs of objects with overlapping VPN domains, only if the objects (that represent VPN sites) are included in the same VPN community. This parameter is also used, if the same destination IP can be reached through more than one VPN community.
<code>traditional</code>	Default parameter. Shows all pairs of objects with overlapping VPN domains.

Example

```
# vpn overlap_encdom communities
The objects Paris and London have overlapping encryption domains.
The overlapping domain is:
10.8.8.1 - 10.8.8.1
10.10.8.0 - 10.10.9.255
- This overlapping encryption domain generates a multiple entry points configuration in
MyIntranet and RemoteAccess communities.
- Same destination address can be reached in more than one community (Meshed, Star). This
configuration is not supported.

The objects Paris and Chicago have overlapping encryption domains. The overlapping domain is:
10.8.8.1 - 10.8.8.1
- Same destination address can be reached in more than one community (MyIntranet, NewStar). This
configuration is not supported.

The objects Washington and Tokyo have overlapping encryption domains.
The overlapping domain is:
10.12.10.68 - 10.12.10.68
10.12.12.0 - 10.12.12.127
10.12.14.0 - 10.12.14.255
- This overlapping encryption domain generates a multiple entry points configuration in Meshed,
Star and NewStar communities.
```

vpn rim_cleanup

Description

Cleans RIM routes.

Syntax

```
vpn rim_cleanup
```

vpn rll

Description

Controls the VPN Route Lookup Layer:

- Saves the Route Lookup Layer information to the specified file.
- Synchronizes the routing table.

Syntax

```
vpn rll
    dump <File>
    sync
```

Parameters

Parameter	Description
dump <File>	Saves the Route Lookup Layer information to the specified file: ▪ ISP Redundancy Default Routes (Next Hop, Interface, Metric) ▪ Route Shadow (Interface and Metric, IP/Mask, Next Hop) ▪ Monitored IP Addresses (Data, IP/Mask)
sync	Synchronizes the routing table.

vpn set_slim_server

Description

This command is deprecated.

Delete the `$FWDIR/conf/slim.conf` file and use the Management Server to configure SSL Network Extender.

As long as the `$FWDIR/conf/slim.conf` file exists, it **overrides** the settings you configure on the Management Server.

vpn set_snx_encdom_groups

Description

Controls the encryption domain per usergroup feature for SSL Network Extender.

Syntax

```
vpn set_snx_encdom_groups
    off
    on
```

Parameters

Parameter	Description
off	Disables the encryption domain per usergroup feature.
on	Enables the encryption domain per usergroup feature.

vpn set_trac

Description

Controls the TRAC server.

Syntax

```
vpn set_trac  
    disable  
    enable
```

Parameters

Parameter	Description
disable	Disables the TRAC server.
enable	Enables the TRAC server.

Example

```
[Expert@MyGW:0]# vpn set_trac enable  
Trac client enabled, Install Policy for this change to take effect  
[Expert@MyGW:0]#  
  
[Expert@MyGW:0]# vpn set_trac disable  
Trac client disabled, Install Policy for this change to take effect  
[Expert@MyGW:0]#
```

vpn shell

Description

VPN Command Line Interface.

Syntax for IPv4

```
vpn shell
```

Syntax for IPv6

```
vpn6 shell
```

Menu Options

```
[Expert@MyGW:0]# vpn shell
?                - This help
..               - Go up one level
quit             - Quit
[interface      ] - Manipulate tunnel interfaces
[show           ] - Show internal data
[tunnels        ] - Manipulate tunnel data
[license        ] - Display SCM licenses
VPN shell:[/] >
```

Menu Sub-Options

```

interface
  add
  modify
  delete
  show
show
  interface
  tunnels
    IKE
      all
      peer <Internal Peer IP>
    IPsec
      all
      peer <Internal Peer IP>
tunnels
  show
    IKE
      all
      peer <Internal Peer IP>
    IPsec
      all
      peer <Internal Peer IP>
  delete
    IKE
      peer <Security Gateway>
      user <Username>
      all
    IPsec
      peer <Security Gateway>
      user <Username>
      all
    all
      IKE
      IPsec
license
  scm
      status
      list

```

Description of Options and Sub-Options

Option	Description
?	Shows the available advanced commands in the current menu level.
..	Goes up one level in the menu.
quit	Quits the VPN shell (available only in the main level).
interface	These commands are deprecated on Gaia OS. Use the applicable options in Gaia Portal or the applicable commands in Gaia Clish. See the R81.10 Gaia Administration Guide.
show	Shows internal data. The available options are: ■ Show and configure tunnel interfaces: <pre>show > interface</pre> These commands are deprecated on Gaia OS. Use the applicable options in Gaia Portal or the applicable commands in Gaia Clish. See the R81.10 Gaia Administration Guide.

Option	Description
	■ Show Security Associations (SAs): show > tunnels The available sub-options are: • Show all IKE SAs show > tunnels > IKE > all Note - This sub-option is the same as: ◦ In the main "vpn tu" on page 252 menu, the option (3) List all IKE SAs for a given peer (GW). ◦ The "vpn tu [-w] list ike" command (see "vpn tu list" on page 257). • Show all IKE SAs for a specified VPN peer: show > tunnels > IKE > peer <Internal Peer IP> Note - This sub-option is the same as: ◦ In the main "vpn tu" on page 252 menu, the option (1) List all IKE SAs. ◦ The "vpn tu [-w] list peer_ike <IP Address>" command (see "vpn tu list" on page 257). • Show all IPsec SAs show > tunnels > IPsec > all Note - This sub-option is the same as: ◦ In the main "vpn tu" on page 252 menu, the option (2) List all IPsec SAs. ◦ The "vpn tu [-w] list ipsec" command (see "vpn tu list" on page 257). • Show all IPsec SAs for a specified VPN peer: show > tunnels > IPsec > peer <Internal Peer IP> Note - This sub-option is the same as: ◦ In the main "vpn tu" on page 252 menu, the option (4) List all IPsec SAs for a given peer (GW). ◦ The "vpn tu [-w] list peer_ipsec <IP Address>" command (see "vpn tu list" on page 257).

Option	Description
tunnels	Shows and deletes Security Associations (SAs). The available options are: ▪ Show Security Associations (SAs): tunnels > show The available sub-options are: • Show all IKE SAs: tunnels > show > IKE > all Note - This sub-option is the same as: ◦ In the main "vpn tu" on page 252 menu, the option (1) List all IKE SAs. ◦ The "vpn tu [-w] list ike" command (see "vpn tu list" on page 257). • Show all IKE SAs for a specified VPN peer: tunnels > show > IKE > peer <Internal Peer IP> Note - This sub-option is the same as: ◦ In the main "vpn tu" on page 252 menu, the option (3) List all IKE SAs for a given peer (GW). ◦ The "vpn tu [-w] list peer_ike <IP Address>" command (see "vpn tu list" on page 257). • Show all IPsec SAs: tunnels > show > IPsec > all Note - This sub-option is the same as: ◦ In the main "vpn tu" on page 252 menu, the option (2) List all IPsec SAs. ◦ The "vpn tu [-w] list ipsec" command (see "vpn tu list" on page 257). • Show all IPsec SAs for a specified VPN peer: tunnels > show > IPsec > peer <Internal Peer IP> Note - This sub-option is the same as: ◦ In the main "vpn tu" on page 252 menu, the option (4) List all IPsec SAs for a given peer (GW). ◦ The "vpn tu [-w] list peer_ipsec <IP Address>" command (see "vpn tu list" on page 257).

Option	Description
	■ Delete Security Associations (SAs): <pre>tunnels > delete</pre> The available sub-options are: • Delete all IKE for a specified VPN peer: <pre>tunnels > delete > IKE > peer <Internal Peer IP></pre> • Delete all IKE for a specified user: <pre>tunnels > delete > IKE > user <Username></pre> • Delete all IKE SAs for all VPN peers and users: <pre>tunnels > delete > IKE > all</pre> <pre>tunnels > delete > all > IKE</pre> • Delete all IPsec SAs for a specified VPN peer: <pre>tunnels > delete > IPsec > peer <Internal Peer IP></pre> Note - This sub-option is the same as: ◦ In the main "vpn tu" on page 252 menu, the option (5) Delete all IPsec SAs for a given peer (GW). ◦ The "vpn tu [-w] del ipsec <IP Address>" command (see "vpn tu del" on page 254). • Delete all IPsec SAs for a specified user: <pre>tunnels > delete > IPsec > user <Username></pre> Note - This sub-option is the same as: ◦ In the main "vpn tu" on page 252 menu, the option (6) Delete all IPsec SAs for a given User (Client). ◦ The "vpn tu [-w] del ipsec <IP Address> <Username>" command (see "vpn tu del" on page 254). • Delete all IPsec SAs for all VPN peers and users: <pre>tunnels > delete > IPsec > all</pre> <pre>tunnels > delete > all > IPsec</pre> Note - This sub-option is the same as: ◦ In the main "vpn tu" on page 252 menu, the option (9) Delete all IPsec SAs for ALL peers and users. ◦ The "vpn tu [-w] del ipsec all" command (see "vpn tu del" on page 254).

Option	Description
license	Shows the SecureClient Mobile (SCM) licenses. The available sub-options are: ▪ Show the current status of SCM licenses: <pre>license > scm > status</pre> ▪ Show the list of SCM licensed devices: <pre>license > scm > list</pre>

vpn show_tcpt

Description

Shows users connected in Visitor Mode.

Syntax

```
vpn show_tcpt
```

vpn sw_topology

 **Note** - R81.10 does not support UTM-1 Edge and Safe@Office devices. The information about this command is provided only to describe the existing syntax option until it is removed completely.

Description

Downloads the topology for a UTM-1 Edge or Safe@Office device.

Syntax

```
vpn [-d] sw_topology -dir <directory> -name <name> -profile
<profile> [-filename <filename>]
```

Parameters

Parameter	Description
-d	Runs the command in debug mode. Use only if you troubleshoot the command itself.  Best Practice - If you use this parameter, then redirect the output to a file, or use the script command to save the entire CLI session.
-dir <directory>	Output directory for file.
-name <name>	Nickname of site, which appears in remote client.
-profile <profile>	Name of the UTM-1 Edge or Safe@Office profile, for which the topology is created.
-filename <filename>	Name of the output file.

vpn tu

Description

Launches the `TunnelUtil` tool, which is used to control VPN tunnels.

General Syntax

```
vpn tu
```

```
vpn tunnelutil
```

Menu Options

```
[Expert@MyGW:0]# vpn tu
```

```
*****          Select Option          *****

(1)          List all IKE SAs
(2)          * List all IPsec SAs
(3)          List all IKE SAs for a given peer (GW)
(4)          * List all IPsec SAs for a given peer (GW)
(5)          Delete all IPsec SAs for a given peer (GW)
(6)          Delete all IPsec SAs for a given User (Client)
(7)          Delete all IPsec+IKE SAs for a given peer (GW)
(8)          Delete all IPsec+IKE SAs for a given User
(Client)
(9)          Delete all IPsec SAs for ALL peers
(0)          Delete all IPsec+IKE SAs for ALL peers

* To list data for a specific CoreXL instance, append "-i
<instance number>" to your selection.

(Q)          Quit

*****
```



Note - When you view Security Associations for a specific VPN peer, you must specify the IP address in dotted decimal notation.

Advanced Syntax

```
vpn tu [-w]
      help
      conn <options>
      del <options>
      list <options>
      mstats
      tlist <options>
```

Parameters

Parameter	Description
help	Shows the available advanced commands.
-w	Shows various warnings on the screen.
conn <options>	Shows VPN information for the specified traffic connections.
del <options>	Deletes IPsec and IKE SAs. See "vpn tu del" on page 254 .
list <options>	Shows IPsec and IKE SAs. See "vpn tu list" on page 257 .
mstats	Shows distribution of VPN tunnels (SPIs) between CoreXL Firewall instances. See "vpn tu mstats" on page 259 .
tlist <options>	Shows information about VPN tunnels. See "vpn tu tlist" on page 260 .

vpn tu del

Description

Deletes IPsec Security Associations (SAs) and IKE Security Associations (SAs).

Syntax for IPv4

```
vpn tu [-w] del
    all
    ipsec
        all
        <IPv4 Address>
        <IPv4 Address> <Username>
    <IPv4 Address>
    <IPv4 Address> <Username>
```

Syntax for IPv6

```
vpn tu [-w] del
    all
    ipsec
        all
        <IPv6 Address>
    <IPv6 Address>
    <IPv6 Address> <Username>
```

Parameters

Parameter	Description
-w	Shows various warnings on the screen.
all	Deletes all IPsec SAs and IKE SAs for all VPN peers and users.  Note - This command is the same as: ▪ In the main "vpn tu" on page 252 menu, the option (0) Delete all IPsec+IKE SAs for ALL peers and users. ▪ In the "vpn shell" on page 243 menu, the option tunnels > delete > all > IKE and the option tunnels > delete > all > IPsec..

Parameter	Description
<code>ipsec <options></code>	Deletes the specified IPsec SAs. The available <options> are: - Delete all IPsec SAs for all peers and users: <code>vpn tu [-w] del ipsec all</code> Note - This command is the same as: - In the main "vpn tu" on page 252 menu, the option (9) Delete all IPsec SAs for ALL peers and users. - In the "vpn shell" on page 243 menu, the option tunnels > delete > all > IPsec. - Delete all IPsec SAs for the specified VPN peer: <code>vpn tu [-w] del ipsec <IP Address></code> Note - This command is the same as: - In the main "vpn tu" on page 252 menu, the option (5) Delete all IPsec SAs for a given peer (GW). - In the "vpn shell" on page 243 menu, the option tunnels > delete > IPsec > peer <Internal Peer IP>. - Delete all IPsec SAs for the specified VPN peer and the specified user: <code>vpn tu [-w] del ipsec <IPv4 Address> <Username></code> Notes: - This command is the same as: - In the main "vpn tu" on page 252 menu, the option (6) Delete all IPsec SAs for a given User (Client). - In the "vpn shell" on page 243 menu, the option tunnels > delete > IPsec > user <Username>. - This command does not support IPv6 addresses.
<code><IP Address></code>	Deletes all IPsec SAs and IKE SAs for the specified VPN peer. Note - This command is the same as the option (7) Delete all IPsec+IKE SAs for a given peer (GW) in the main "vpn tu" on page 252 menu.

Parameter	Description
<i><IP Address></i> <i><Username></i>	Deletes all IPsec SAs and IKE SAs for the specified VPN peer and the specified user.  Note - This command is the same as the option (8) Delete all IPsec+IKE SAs for a given User (Client) in the main <i>"vpn tu" on page 252</i> menu.

vpn tu list

Description

Shows IPsec SAs and IKE SAs.

Syntax for IPv4 and IPv6

```
vpn tu [-w] list
    ike
    ipsec
    peer_ike <IP Address>
    peer_ipsec <IP Address>
    tunnels
```

Parameters

Parameter	Description
-w	Shows various warnings on the screen.
ike	Shows all IKE SAs.  Note - This command is the same as: ■ In the main <i>"vpn tu" on page 252</i> menu, the option (1) List all IKE SAs. ■ In the <i>"vpn shell" on page 243</i> menu, the option show > tunnels > IKE > all or the option tunnels > show > IKE > all.
ipsec	Shows all IPsec SAs.  Note - This command is the same as: ■ In the main <i>"vpn tu" on page 252</i> menu, the option (2) List all IPsec SAs. ■ In the <i>"vpn shell" on page 243</i> menu, the option show > tunnels > IPsec > all or the option tunnels > show > IPsec > all.

Parameter	Description
<code>peer_ike <IP Address></code>	Shows all IKE SAs for the specified VPN peer.  Note - This command is the same as: ■ In the main <i>"vpn tu" on page 252</i> menu, the option (3) List all IKE SAs for a given peer (GW). ■ In the <i>"vpn shell" on page 243</i> menu, the option <code>show > tunnels > IKE > peer <Internal Peer IP></code> or the option <code>tunnels > show > IKE > peer <Internal Peer IP></code>.
<code>peer_ipsec <IP Address></code>	Shows all IPsec SAs for the specified VPN peer.  Note - This command is the same as: ■ In the main <i>"vpn tu" on page 252</i> menu, the option (4) List all IPsec SAs for a given peer (GW). ■ In the <i>"vpn shell" on page 243</i> menu, the option <code>show > tunnels > IPsec > peer <Internal Peer IP></code> or the option <code>tunnels > show > IPsec > peer <Internal Peer IP></code>.
<code>tunnels</code>	Shows information about VPN tunnels. In addition, see the <i>"vpn tu tlist" on page 260</i> command.

vpn tu mstats

Description

Shows the distribution of VPN traffic between CoreXL Firewall instances.

For more information, see [sk118097 - MultiCore Support for IPsec VPN in R80.10 and above](#).

Syntax for IPv4

```
vpn tu [-w] mstats
```

Syntax for IPv6

```
vpn6 tu [-w] mstats
```

Parameters

Item	Description
-w	Shows various warnings on the screen.

Example for IPv4

```
[Expert@MyGW:0]# vpn tu mstats

Instance# # of inSPIs # of outSPIs
0 182 170
1 184 176
2 191 174
3 215 197
4 237 227
5 191 176
6 180 170
7 190 166
8 171 160
9 199 187
-----
Summary: 1940 1803

[Expert@MyGW:0]#
```

Example for IPv6

```
[Expert@MyGW:0]# vpn6 tu mstats

Instance# # of inSPIs # of outSPIs
0 238 228
1 224 214
-----
Summary: 462 442

[Expert@MyGW:0]#
```

vpn tu tlist

Description

Shows information about VPN tunnels.

Syntax for IPv4

```
vpn tu [-w] tlist  
        {-h | -help}  
        [clear]  
        [start]  
        [state]  
        [stop]  
        [<Sort Options>]
```

Syntax for IPv6

```
vpn6 tu [-w] tlist  
        {-h | -help}  
        [clear]  
        [start]  
        [state]  
        [stop]  
        [<Sort Options>]
```

Parameters

Parameter	Description
-w	Shows various warnings on the screen.
-h -help	Shows the built-in usage.
clear	Clears the Tunnel List volume statistics.
start	Turns on the Tunnel List volume statistics.
state	Shows the current Tunnel List volume statistics state.
stop	Turns off the Tunnel List volume statistics.

Parameter	Description
<Sort Options>	The available sort options are: ▪ -b - Sorts by total (encrypted + decrypted) bytes. ▪ -d - Sorts by inbound (decrypted) bytes. ▪ -e - Sorts by outbound (encrypted) bytes. ▪ -i - Combines list rows for each CoreXL Firewall instance with accumulated traffic. Default order is descending by total bytes. ▪ -m - Sorts by MSPI. ▪ -n - Sorts by VPN peer name. ▪ -p <IP Address> - Shows tunnels only for a VPN peer with the specified IP address. ▪ -r - Sorts in reverse order. ▪ -s - Sorts by SPI. ▪ -t - Combines list rows for each VPN peer with accumulated traffic. Default order is descending by total bytes. ▪ -v - Verbose mode, prints a header message for each option. If you specify more than one sort option, you can: ▪ Separate the options with spaces: <pre>... -<option1> -<option2> -<option3></pre> For example: -v -t -b -r ▪ Write the options together: <pre>... -<option1><option2><option3></pre> For example: -vtbr

Example for IPv4

```
[Expert@MyGW:0]# vpn tu tlist
```

```
+-----+-----+-----+
| Peer: 172.29.7.134 (b61cef72a222a909) | MSA: ffffc20020e34530 | i: 2 ref: 11 |
| Methods: ESP Tunnel AES-128 SHA1 | | i: 5 ref: 2 |
| My TS: 0.0.0.0/0 | | |
| Peer TS: 172.29.7.134 | | |
| User: user3 | | |
| MSPI: b7 (i: 5) | Out SPI: c95d172c | |
+-----+-----+-----+
[Expert@MyGW:0]#
```

vpn ver

Description

Shows the major version number and build number of the VPN kernel module.

Syntax

```
vpn ver [-k] [-f <filename>]
```

Parameters

Parameter	Description
-k	Shows the version name and build number and the kernel build number.
-f	Saves the information to the specified text file.

Example

```
[Expert@MyGW:0]# vpn ver -k  
This is Check Point VPN-1(TM) R81.10 - Build 123  
kernel: R81.10 - Build 456  
[Expert@MyGW:0]#
```

mcc

Description

The VPN Multi-Certificate CA (MCC) commands let you manage certificates and Certificate Authorities on a Security Management Server or Domain Management Server:

- Shows Certificate Authorities
- Shows certificates
- Adds certificates
- Deletes certificates

Important:

- Before you run this command, you must close all SmartConsole clients, Database Tool (GuiDBEdit Tool) clients (see [sk13009](#)), and "dbedit" clients (see [sk13301](#)) to prevent a lock of the management database. The only exceptions are the "mcc lca" and "mcc show" commands.
- The `mcc` commands require the `cpca` process to be up and running. Run this command:

```
ps auxw | egrep "cpca|COMMAND"
```

- On a Multi-Domain Server, you must run this command in the context of the applicable Domain Management Server:

```
mdsenv <IP Address or Name of Domain Management Server>
```

Syntax

```
mcc
    -h
    add <options>
    add2main <options>
    del <options>
    lca
    main2add <options>
    show <options>
```

Parameters

Parameter	Description
-h	Shows the built-in usage.

Parameter	Description
<code>add <options></code>	Adds certificates. See "mcc add" on page 265 .
<code>add2main <options></code>	Promotes an additional certificate to be the main certificate. See "mcc add2main" on page 266 .
<code>del <options></code>	Deletes certificates. See "mcc del" on page 267 .
<code>lca</code>	Shows Certificate Authorities. See "mcc lca" on page 268 .
<code>main2add <options></code>	Adds main certificate to additional certificates. See "mcc main2add" on page 269 .
<code>show <options></code>	Shows certificates. See "mcc show" on page 270 .

mcc add

Description

Adds a certificate stored in DER format in a specified file, as an additional certificate to the specified CA. The new certificate receives an index number higher by one than the highest existing certificate index number.

The new certificate receives an index number higher by one than the highest existing certificate index number.

Syntax

```
mcc add <CA Name> <Certificate File>
```



Note

On a Multi-Domain Server, you must run this command in the context of the applicable Domain Management Server:

```
mdsend <IP Address or Name of Domain Management Server>
```



Important - Before you run this command, you must close all SmartConsole clients, Database Tool (GuiDBEdit Tool) clients (see [sk13009](#)), and "dbedit" clients (see [sk13301](#)) to prevent a lock of the management database.

Parameters

Parameter	Description
<i><CA Name></i>	Specifies the name of the CA, as defined in the Management Server database.
<i><Certificate File></i>	Specifies the path and the name of the certificate file. To show the main certificate of a CA, omit this parameter.

Example - Add the certificate stored in the `/var/log/Mycert.cer` file to the CA called "MyCA"

```
mcc add MyCA /var/log/Mycert.cer
```

mcc add2main

Description

Copies the additional certificate of the specified index number of the specified CA to the main position and overwrites the previous main certificate.

Syntax

```
mcc add2main <CA Name> <Certificate Index Number>
```



Note

On a Multi-Domain Server, you must run this command in the context of the applicable Domain Management Server:

```
mdsenv <IP Address or Name of Domain Management Server>
```



Important - Before you run this command, you must close all SmartConsole clients, Database Tool (GuiDBEdit Tool) clients (see [sk13009](#)), and "dbedit" clients (see [sk13301](#)) to prevent a lock of the management database.

Parameters

Parameter	Description
<i><CA Name></i>	Specifies the name of the CA, as defined in the Management Server database.
<i><Certificate Index Number></i>	Specifies the certificate index number.

Example - Copy certificate #1 of a CA called "MyCA" to the main position

```
mcc add2main MyCA 1
```

mcc del

Description

Removes the additional certificate of the specified index number from the specified CA.

Greater index numbers (of other additional certificates) are reduced by one.

Syntax

```
mcc del <CA Name> <Certificate Index Number>
```



Note

On a Multi-Domain Server, you must run this command in the context of the applicable Domain Management Server:

```
mdsenv <IP Address or Name of Domain Management Server>
```



Important - Before you run this command, you must close all SmartConsole clients, Database Tool (GuiDBEdit Tool) clients (see [sk13009](#)), and "dbedit" clients (see [sk13301](#)) to prevent a lock of the management database.

Parameters

Parameter	Description
<CA Name>	Specifies the name of the CA, as defined in the Management Server database.
<Certificate Index Number>	Specifies the certificate index number.

Example - Remove certificate #1 of a CA called "MyCA"

```
mcc del MyCA 1
```

mcc lca

Description

Shows all Certificate Authorities (CAs) defined in the Management Server database, with the number of additional CA certificates for each CA.

Syntax

```
mcc lca
```



Note

On a Multi-Domain Server, you must run this command in the context of the applicable Domain Management Server:

```
mdsensv <IP Address or Name of Domain Management Server>
```

Example

```
[Expert@MGMT:0]# mcc lca
MCC: Here is a list of the CAs, with the number of additional CA certificates
      1. internal_ca (0)
[Expert@MGMT:0]#
```

mcc main2add

Description

Copies the main certificate of the specified CA to an additional position.

The copied certificate receives an index number higher by one than the highest existing certificate index number.

Syntax

```
mcc main2add <CA Name>
```



Note

On a Multi-Domain Server, you must run this command in the context of the applicable Domain Management Server:

```
mdsenv <IP Address or Name of Domain Management Server>
```



Important - Before you run this command, you must close all SmartConsole clients, Database Tool (GuiDBEdit Tool) clients (see [sk13009](#)), and "dbedit" clients (see [sk13301](#)) to prevent a lock of the management database.

Parameters

Parameter	Description
<CA Name>	Specifies the name of the CA, as defined in the Management Server database.

Example

The CA called "MyCA" has a main certificate and one additional certificate.

If you run this command, then the CA will have two additional certificates, and additional certificate #2 will be identical to the main certificate:

```
mcc main2add MyCA
```

mcc show

Description

Shows details for a specified certificate of a specified CA.

Syntax

```
mcc show <CA Name> [<Certificate Index Number>]
```



Note

On a Multi-Domain Server, you must run this command in the context of the applicable Domain Management Server:

```
mdserv <IP Address or Name of Domain Management Server>
```

Parameters

Parameter	Description
<CA Name>	Specifies the name of the CA, as defined in the Management Server database.
<Certificate Index Number>	Optional. Specifies the certificate index number. To show the main certificate of a CA, omit this parameter.

Example 1 - Show certificate #1 of a CA called MyCA

```
mcc show MyCA 1
```

Example 2 - Show certificate of a CA called "internal_ca"

```
[Expert@MGMT:0]# mcc lca
MCC: Here is a list of the CAs, with the number of additional CA certificates
1. internal_ca (0)
[Expert@MGMT:0]#

[Expert@MGMT:0]# mcc show internal_ca
PubKey:
Modulus:
ae b3 75 36 64 e4 1a 40 fe c2 ad 2f 9b 83 0b 45 f1 00 04 bc
3f 77 77 76 d1 de 8a cf 9f 32 78 8b d4 b1 b4 be db 75 cc c8
... ..
a3 9d 8b 0a de 05 fb 5c 44 2e 29 e3 3e f4 dd 50 01 0f 86 9d
55 16 a3 4d f8 90 2d 13 c6 c1 28 57 f8 3e 7c 59
Exponent: 65537 (0x10001)

X509 Certificate Version 3
refCount: 1
Serial Number: 1
Issuer: O=MyServer.checkpoint.com.s6t98x
Subject: O=MyServer.checkpoint.com.s6t98x
Not valid before: Sun Apr 8 13:41:00 2018 Local Time
Not valid after: Fri Jan 1 05:14:07 2038 Local Time
Signature Algorithm: RSA with SHA-256 Public key: RSA (2048 bits)
Extensions:
    Key Usage:
        digitalSignature
        keyCertSign
        cRLSign
    Basic Constraint (Critical):
        is CA

[Expert@MGMT:0]#
```

Working with Kernel Parameters

See the [R81.10 Quantum Security Gateway Guide](#) > Chapter "Working with Kernel Parameters".

Kernel Debug

See the [R81.10 Quantum Security Gateway Guide](#) > Chapter "Kernel Debug on Security Gateway".

Appendix

Configuring specific settings for each VPN Community

By default, many global VPN settings you configure in SmartConsole (in **Global properties**) apply to all managed Security Gateways.

You can override these global settings for a specific VPN Community:

- Time interval, in seconds, for sending life sign packets.
- Maximum number of concurrent Internet Key Exchange (IKE) negotiations that occur at the same time.

Procedure

1. Back up the Security Management Server / applicable Domain Management Server.

Refer to:

- [sk108902 - Best Practices - Backup on Gaia OS.](#)
- [sk91400 - System Backup and Restore feature in Gaia.](#)
- [sk98153 - How to take a snapshot of Endpoint Security Management Server database.](#)

2. Close **all** SmartConsole windows.

Note - To make sure there are no active sessions, run the "`cpstat mg`" command in the Expert mode on the Security Management Server / in the context of *each* Domain Management Server.

3. Connect with the Database Tool (GuiDBEdit Tool) to the Security Management Server / applicable Domain Management Server.
4. In the upper left pane, click **Tables > Managed Objects > Communities**.
5. In the upper right pane, select the applicable VPN Community.
6. Press **CTRL+F** (or click **Search > Find**).
7. Enter the name of the applicable attribute
See the summary table below.
8. Click **Find Next**.
9. In the lower pane, right click the attribute name.

10. Select **Edit**.
11. Configure the applicable value.
12. Click **OK**.
13. Save the changes. Click **File > Save All**.
14. Close the Database Tool (GuiDBEdit Tool).
15. Connect with SmartConsole to the Security Management Server or the applicable Domain Management Server.
16. Install the Access Control policy on all Security Gateways and Cluster objects that participate in this VPN Community.

Attributes

VPN Feature	Attribute	Valid Values	Description
Dead Peer Detection (DPD)	life_sign_timeout	Range: 5-3600 sec Default: 40 sec	Controls the DPD timeout in a VPN community
	life_sign_transmitter_interval	Range: 5-600 sec Default: 10 sec	Controls the DPD transmission interval in a VPN community
	override_global_settings_for_life_sign_intervals	Range: true, false Default: false	Override the global settings for life sign intervals

VPN Feature	Attribute	Valid Values	Description
IKE Negotiations	max_negotiations	Range: 1-10000 Default: 10000	Controls the number of IKE negotiations in a VPN community This helps VPN Gateways to cope with a situation of boot-storm over slow WAN links. After a new IKE/IPsec (IKEv1 or IKEv2) negotiation starts with a VPN peer, the VPN Gateway allows or denies it based on the configured threshold. Note- IKE informational packets (for example, DPD) are not counted as negotiation.
	override_max_concurrent_ike_negotiation	Range: true, false Default: false	Overrides the Maximum number of concurrent IKE negotiations per community