

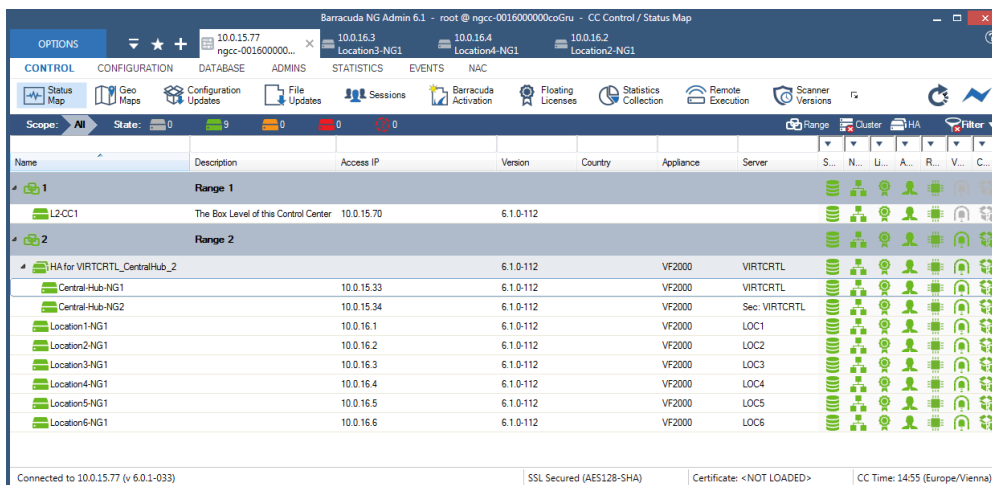
Implementation Guide - VPN Network with Static Routing

<https://campus.barracuda.com/doc/53248635/>

This guide contains advanced topics and concepts. Follow the links in each section for step-by-step instructions on how to configure the following features.

- [Site-to-Site VPN](#)
 - [TINA Tunnel Settings](#)
 - [Traffic Intelligence](#)
 - [GTI Editor](#)
 - [\(optional\) WAN Optimization](#)
- [Forwarding Firewall](#)
 - [Connection Objects](#)
 - [Traffic Shaping](#)
- [Distributed Firewall / Shared Services](#)
- [Central Management](#)
- [Global Firewall Objects](#)
- [Repositories](#)


The Barracuda NextGen Firewall F-Series and NextGen Control Center are both designed for the quick deployment and easy management of a large number of F-Series Firewalls. The Control Center offers features that allow you to apply the parts of the configuration that are the same or similar on all the NextGen Firewall F-Series units. In this example, we are configuring a VPN hub with remote firewalls connected via Site-to-Site VPN. By employing repositories, global/range/cluster Firewall Objects, and shared services, the configuration path is designed to be as efficient as possible. Employing Firewall Objects enables you to quickly change or add additional networks without having to change the configuration of your VPN network.



The screenshot shows the Barracuda NG Admin 6.1 interface. The top navigation bar includes tabs for CONTROL, CONFIGURATION, DATABASE, ADMINS, STATISTICS, EVENTS, and NAC. Below this is a toolbar with various icons for status, maps, configuration, updates, sessions, activation, licenses, statistics, remote execution, and scanner versions. The main content area displays a 'Status Map' with a table of firewall objects.

Name	Description	Access IP	Version	Country	Appliance	Server	S...	N...	U...	A...	R...	V...	C...
1	Range 1												
L2CC1	The Box Level of this Control Center	10.0.15.70	6.1.0-112										
2	Range 2												
HA for VIRTCTRL_CentralHub_2			6.1.0-112		VF2000	VIRTCTRL							
Central-Hub-NG1		10.0.15.33	6.1.0-112		VF2000	VIRTCTRL							
Central-Hub-NG2		10.0.15.34	6.1.0-112		VF2000	Sec: VIRTCTRL							
Location1-NG1		10.0.16.1	6.1.0-112		VF2000	LOC1							
Location2-NG1		10.0.16.2	6.1.0-112		VF2000	LOC2							
Location3-NG1		10.0.16.3	6.1.0-112		VF2000	LOC3							
Location4-NG1		10.0.16.4	6.1.0-112		VF2000	LOC4							
Location5-NG1		10.0.16.5	6.1.0-112		VF2000	LOC5							
Location6-NG1		10.0.16.6	6.1.0-112		VF2000	LOC6							

Connected to 10.0.15.77 (v 6.0.1-033) | SSL Secured (AES128-SHA) | Certificate: <NOT LOADED> | CC Time: 14:55 (Europe/Vienna)




Section 1 - Preparation

Before You Begin

- NextGen Control Center must use the same or later firmware version as the managed unit.
- For future reference, create a table with the following information for each NextGen Firewall F-Series:
 - A list of local VPN networks to route through the tunnels.
 - All public IP addresses to be used for the VPN service.

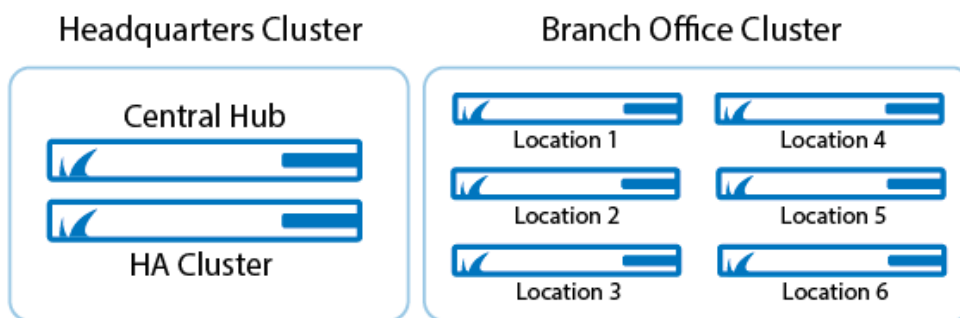
Information for all Firewalls in our Network

NextGen Firewall F-Series	Local VPN Networks	MIP	Public IP Addresses
Hub/HQ	10.0.15.0/24, 172.16.0.0/24	10.0.15.33 and 10.0.15.88	62.99.0.XX, 194.93.0.XX, 10.20.0.7 (MPLS)
Location 1	10.0.40.0/24, 172.16.1.0/24	10.0.40.1	212.86.0.XX, 80.130.45.XX, 10.21.0.7 (MPLS)
Location 2	10.0.51.0/24, 172.16.2.0/24	10.0.51.1	213.47.0.XX, 10.22.0.7 (MPLS)
Location 3	10.0.60.0/24, 172.16.3.0/24	10.0.60.1	214.51.2.80
Location 4	10.0.70.0/24, 10.0.71.0/24, 10.0.72.0/24	10.0.70.1	dynamic
Location 5	10.0.80.0/24, 172.16.5.0/24	10.0.80.1	dynamic
Location 6	10.0.81.0/24, 172.16.6.0/24	10.0.81.1	dynamic



Preparing the NextGen Control Center

Organizing your Firewalls into Clusters and Ranges



To be able to use distributed services and global firewall objects efficiently, you must organize your firewalls in the respective clusters and ranges. If all your firewalls are running the same firmware version, you can use just one cluster. If some of your firewalls are running an older firmware version, you will need to create a cluster for each version. In this example, the headquarters and branch office firewalls will each use their own cluster since using multiple clusters makes managing the configuration easier.

- **Branch Office Clusters** – Add all branch office NextGen Firewall F-Series units to this cluster. All units must use the same firmware version. These NextGen Firewall F-Series units will share a Distributed Firewall service.
- **Headquarters Cluster** – The central firewall, which is used as the VPN hub, shares a smaller part of the configuration of the branch office firewall units and does not use the distributed services.

For more information, see [How to Manage Ranges and Clusters](#).

Global Firewall Objects

Global Firewall Objects allow you to enter the network addresses once for all the networks, public IP addresses, and special servers, and then to reuse them when configuring the services. A Global Firewall Object on the global or range level can be overridden by a different IP address or network on the range or cluster level. This allows for one-time configurations in cases where one cluster uses a different IP address or network from all other configurations. You can also employ this functionality to enforce the usage of the same firewall object names for all your configurations. This allows you to create repository entries to be reused for all clusters. Site-specific firewall objects are globally defined in name and type, and the IP addresses or networks are entered in the **Server Properties** of the virtual servers. Site-specific Network Objects can be used only in the Forwarding and Distributed Firewall services.

Create site-specific network objects for networks or server IP addresses (for example, your exchange servers that differ for each location). Avoid using the generic network object type; instead, be as specific as possible. For this example, create the following global network objects for each location:


- **Single IP Address Network Object** – Create a network object for each public IP address used by the NextGen Firewall F-Series. Also create an empty network object for the VPN next hop interface IP address, which will be filled in later.
- **Single Network Addresses** – Create a network object for each network routed through the tunnel.
- **List of Network Addresses** – When using multiple networks in the same location, it is useful to have a network object that references all the networks in that location. This network object is updated automatically whenever one of the network objects it references is updated.

For more information, see [Global Firewall Objects](#).

(optional) Firewall Objects Naming Conventions

Although not required, following a naming convention for the global firewall objects simplifies configuration. It also lets you know at first glance what data is stored in the global firewall object. Also setting the **Network Color** makes a network object easily identifiable in the GUI. For this example, we are using:

<Location><location number>_<Type(NET, VIP, ...)>_<modifier. E.g., ALL, numbering, etc...>

 Firewall Objects

Rule Lists

- Access Rules
- Application Rules
- Object Viewer...

Firewall Objects

- Networks
- Applications
- URL Filter
- Services
- User and Groups
- Connections
- Schedules
- Interface Groups
- Proxy ARPs
- Generic IPS Patterns
































Rule List Verification

Edit Networks

- Edit...
- New...
- Delete
- Set Color...
- Copy
- Paste
- Referenced By...
- Resolve DNS Entries

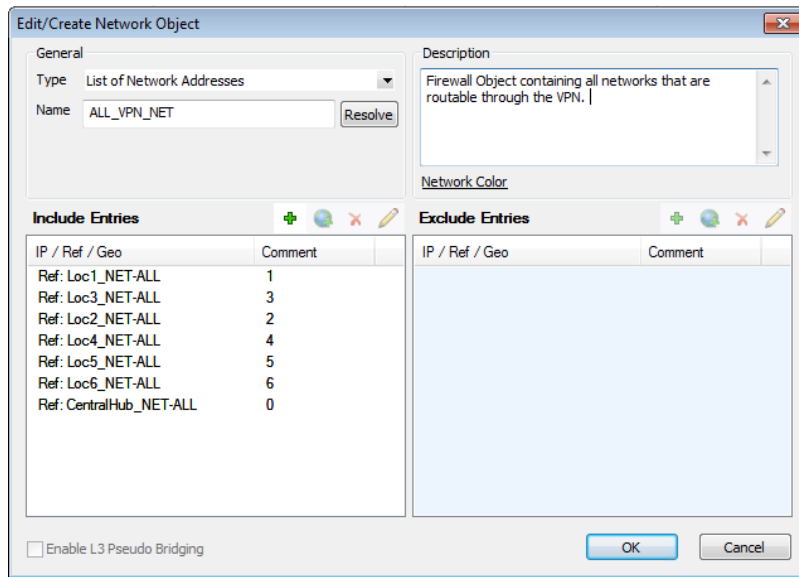
Information

- Setup...
- Export...
- Import...
- Reload Externals

Network objects				
Name	Referenced by	Entries	Comment	
SITE SPECIFIC				
 MailServer	0 (0)		Site-Specific mail server address	
STATIC				
 ALL_VPN_NET	0 (3)	Ref: Loc1_NET-ALL , Ref: ...		
Any	2 (2)	0.0.0.0/0	All IPv4 addresses	
 Central-Hub_NET1	1 (1)	10.0.15.0/24	0	
 Central-Hub_VPNR11	0 (0)	192.168.20.1	0	
 CentralHub-NET2	1 (1)	172.16.1.0/24	0	
 CentralHub_JSP1	0 (7)	62.99.0.80	0	
 CentralHub_JSP2	0 (0)	194.93.0.80	0	
 CentralHub_MPLS_IP	0 (0)	10.20.0.80	0	
 CentralHub_NET-ALL	1 (1)	Ref: Central-Hub_NET1 , R...	0	
 Loc1_JSP1	0 (1)	212.86.0.80	1	
 Loc1_JSP2	0 (1)	80.130.45.80	1	
 Loc1_MIP	0 (0)		1	
 Loc1_MPLS_IP	0 (1)	10.21.0.80	1	
 Loc1_NET-ALL	1 (2)	Ref: Loc1_NET1 , Ref: Loc...	1	
 Loc1_NET1	1 (2)	10.0.40.0/24	1	
 Loc1_NET2	1 (2)	172.16.1.0/24	1	
 Loc1_VIP	0 (1)	10.0.16.1	1	
 Loc1_VPNR11	0 (0)	192.168.20.2	1	
 Loc1_VS_IP1	0 (0)	10.0.40.1	1	
 Loc2_JSP1	0 (1)	213.47.0.80	2	
 Loc2_MIP	0 (0)	10.0.51.1	2	
 Loc2_MPLS_IP	0 (0)	10.22.0.7	2	
 Loc2_NET-ALL	1 (2)	Ref: Loc2_NET1 , Ref: Loc...	2	
 Loc2_NET1	1 (2)	10.0.51.0/24	2	
 Loc2_NET2	1 (2)	172.16.2.0/24	2	
 Loc2_VIP	0 (1)	10.0.16.2	2	
 Loc2_VPNR11	0 (0)	192.168.20.3	2	
 Loc2_VS_IP1	0 (2)	10.0.51.3	2	
 Loc3_JSP1	0 (2)	214.51.2.34	3	
 Loc3_JSP1_BOX	0 (1)	214.51.2.35	3	
 Loc3_MIP	0 (1)	10.0.60.1	3	

Create Global Firewall Objects

Firewall Object Name	Type	Include Entries
RemoteMGMTIPS	List of IP Addresses	Include all IP addresses that must be accessed through the remote management tunnel.
ALL_NET	List of Network Addresses	Contains all networks that are allowed to send and receive traffic through the VPN tunnels. Only add the locX_NET_ALL network object for each location to this network object.
MailServer	Site-Specific Network Object	In this example, the mail server IP address must be configured for each location.



IP / Ref / Geo	Comment
Ref: Loc1_NET-ALL	1
Ref: Loc3_NET-ALL	3
Ref: Loc2_NET-ALL	2
Ref: Loc4_NET-ALL	4
Ref: Loc5_NET-ALL	5
Ref: Loc6_NET-ALL	6
Ref: CentralHub_NET-ALL	0

Create a Repository

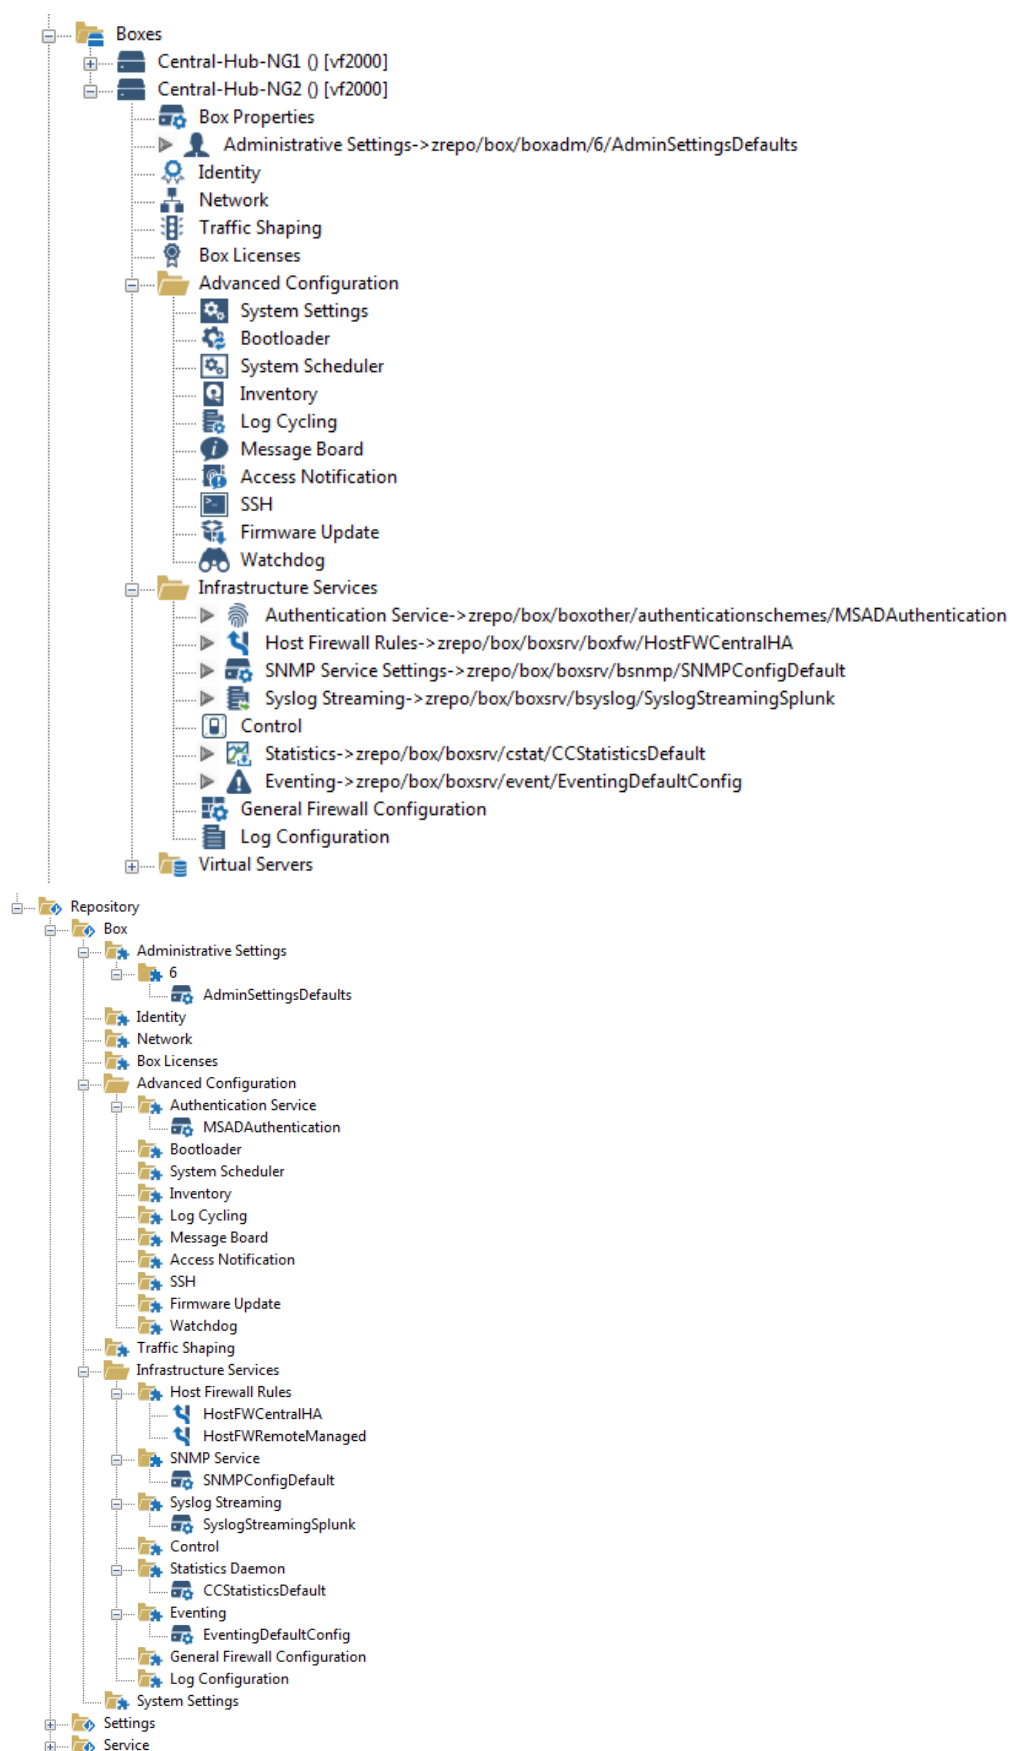
The repository stores preconfigured configuration nodes that can be linked or copied to an individual NextGen Firewall F-Series. You can store several versions of the same configuration node.

- **Creating new Firewalls** – Link repository entries to the **Default Box** of the cluster. All the links will be used when new boxes are created in the cluster. You do not have to set up new firewalls in that cluster from scratch.
- **Existing Firewalls** – Prepare the configuration of a service in the repository. Depending on the amount of customization necessary, link or copy the repository entry. Settings for linked repository entries can be overridden.

The nodes stored and used in the repository depend on the network and personal preference of the admin. The following nodes are frequently used:

- Administrative Settings
- Authentication Service
- Host Firewall Rules
- SNMP Service Settings
- Syslog Streaming
- Statistics
- Eventing

For more information, see [Repositories](#).



Preparing the Managed Firewalls


The NextGen Firewall F-Series units must be managed by the Control Center and have connection to the Internet.

For more information, see [Central Management](#) and [WAN Connections](#).

Create Services

You must create the necessary services for every virtual server:

- **Forwarding Firewall / Distributed Firewall Service** – Reduce the overhead of maintaining a large number of very similar rulesets by using a Distributed Firewall for all firewalls in a cluster.
- **VPN Service**



Section 2 - Setup Overview

Site-to-Site VPN Tunnel Network

TINA VPN Tunnels

Site-to-Site VPN tunnels on the NextGen Firewall F-Series use the Barracuda-proprietary VPN protocol.

TINA offers many enhancements not featured in the standard IPsec protocol, such as Traffic Intelligence, Traffic Compression, and WAN Optimization. Traffic Intelligence allows you to prioritize data flow and distribute VPN traffic between multiple Internet connections. WAN Optimization and Compression reduces the amount of traffic sent through the tunnel by using data deduplication.

Depending on the type of traffic you are sending through the VPN tunnel, choose one of the following VPN transport modes:

- **UDP** – UDP encapsulation benefits from the low-overhead, reduced latency, and NAT traversal capabilities of the UDP protocol. UDP has no error checking, which may be a problem for connections with high packet loss or if VPN traffic largely consists of UDP connections.
- **TCP** – TCP offers transport reliability and NAT traversal capabilities. It is the only available option if you are behind a proxy. If you must connect through an HTTP proxy, port 443 can also be used.
- **Hybrid (TCP & UDP)** – A Hybrid mode tunnel encapsulates TCP in UDP and UDP in TCP to balance the strengths of each protocol with optimal transport reliability. Latency-critical UDP traffic should not be sent in Hybrid mode because the TCP transport mode may increase the latency.
- **ESP** – ESP is the native IPsec protocol, and as a layer 3 protocol, it offers the best performance. NAT traversal is not possible.
- **Routing** – No encapsulation is performed for this transport mode.

Firewall or Distributed Firewall

The Distributed Firewall is designed to lower the overhead of maintaining a large number of firewalls where each firewall services contains the same access rules. To change an access rule for every remote location, you only have to change one access rule in the Global Ruleset of the Distributed Firewall. Unique access rules can still be created in the special or local ruleset that are specific to each location. Since the central location does not share the same access rules, using a forwarding firewall service instead of the distributed firewall service is the better choice.

Section 3 - Configuration Tasks for Every NextGen Firewall F-Series

The following tasks must be completed for each unit:

- Create network objects for network and public IP addresses. For example, **HQ_LAN, HQ_DMZ, HQ_ALL_NET**
- Configure VPN GTI Settings.
- Configure the **GTI Networks** and enter the IP addresses and networks for the site-specific network objects you created in the global firewall objects.


Create Global Firewall Objects

Firewall Object Name	Type	Include Entries
LocX_ISP1	Single IP Address	Public IP address for the first ISP
LocX_ISP2	Single IP Address	Public IP address for the second ISP
LocX_MPLS_IP	Single IP Address	IP address for the MPLS connection
LocX_VIP	Single IP Address	VIP IP address
LocX_VS_IP1	Single IP Address	Internal IP address of the virtual server.
LocX_NET1	Single Network Address	First network in location X
LocX_NET2	Single Network Address	Second network in location X
LocX_NET3	Single Network Address	Third network in location X
LocX_NET_ALL	List of Network Addresses	LocX_NET1, LocX_NET2, LocX_NET3 - all networks used in location X

Add List of Local Network to Global Firewall Object

Add the **Loc<Location Number>_NET_ALL** network object to the **ALL_NETS** network object. We will use this network object later for the access rules.

For example, for Location 1:

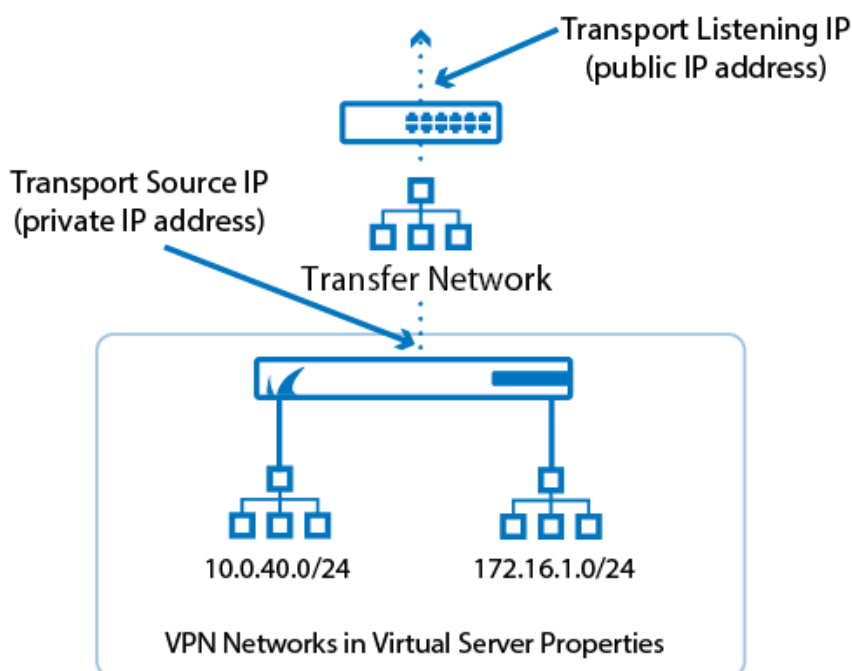


VPN Service Properties

Configure the IP address the VPN service listens on. If you are using firewalls with a dynamic Internet connection (DHCP, xDSL,...), use 127.0.0.1 as the **Service IP** and create an App Redirect access rule to redirect VPN traffic to the VPN service.

For more information, see [How to Configure VPN Access via a Dynamic WAN IP Address](#).

VPN GTI Settings



The GTI Editor uses the public IP addresses in the **VPN GTI Settings** and the **VPN Group** information to create the VPN tunnels. If you want to use all the service IP addresses you set in the **VPN Service Properties**, you can use the default values. You can also explicitly enter the IP address the remote NextGen Firewall F-Series connects to (**Transport Listening IP**) and the Service IP address (**Transport Source IP**) in the **GTI Editor Settings**. These two IP addresses will be the same if your NextGen Firewall F-Series connects directly to the Internet. Setting the IP addresses explicitly is useful when configuring more than two public IP addresses or when safeguarding against breaking your VPN configuration in case the first or second IP address of the virtual server is changed. If you are using only active VPN connections from this VPN service, you can disable the Transport Listening IP by entering 127.0.0.1 as the **Explicit Transport Listening IP**.

Transport Settings for Firewalls with Static Internet Connections:

- **Transport Source IP** - Select **<All-Service-IPs>**.
- **Transport Listening IP** - Select **<Use-Transport-Source-IP>**.

Transport Settings for Firewalls with a Dynamic Internet Connection:

- **Transport Source IP – Dynamic(via-routing).** The NextGen Firewall F-Series uses a routing table lookup to determine the source IP.
- **Transport Listening IP** – Select **Explicit** to not accept incoming VPN connections on dynamic Internet connections. Otherwise, select **<Use-Transport-Source-IP>** if you want to use DynDNS hostnames for incoming VPN connections.
- **Explicit Transport Listening IP** – Enter 127.0.0.1 if you are not going to handle incoming VPN connections.

For more information, see [How to Configure VPN GTI Settings for a VPN Service](#).

Add Networks to the Virtual Server Properties

The networks used as local networks by the GTI Editor are configured in the **Server Properties** of each NextGen Firewall F-Series. Use the **LocX_NET_ALL** network object. Referencing network objects instead of directly entering the networks has the advantage that adding a network to a location is as simple as editing the corresponding network object.

For more information, see [Virtual Servers and Services](#).

Section 4 - VPN GTI Editor

For a network with a large number of Site-to-Site VPN tunnels, it is not practical to configure each tunnel separately on both endpoints. The GTI Editor on the Control Center simplifies and automates this task. Add the VPN services managed on the Control Center into a VPN group. Each VPN group shares VPN configuration settings for encryption, Transport (TCP/UDP/TCP&UDP), and basic Traffic Intelligence configuration.

Create VPN Group

Create the VPN group for this example. Depending on your setup you may need more than one VPN Group to accurately depict your network. Note that using multiple VPN groups will remove the ability to automatically create a fully meshed network. Create the group using the following settings:

- **Transport, Encryption, and Authentication** – Accept the default, or change to match your internal encryption guidelines.
- **Packet Balancing** – Select only **Cycle within a Transport Class** if all of your transports have the same bandwidth and packet round-trip times. If this is not the case, configure session-based balancing in the Connection Object of the access rule that is handling VPN traffic.
- **Meshed** – When set to **yes**, the GTI Editor will automatically create a fully meshed VPN network where all firewalls are connected to each other. Depending on the number of Firewalls involved, this may not be desired because too many site-to-site tunnels can overload the smaller NextGen Firewall F-Series models. In this case, you will have to create the tunnel via drag and

drop. For our example, select **yes** since we want the GTI Editor to create the tunnels.

- **Service Placement** – Classic circular. This setting will automatically arrange the VPN services around the VPN service marked as **Hub**.
- **(optional) WANOptPolicy** – Select the WANOpt policy for this VPN group.

TINA Properties		Edit IPSec
Transport	UDP	
Encryption	AES	
Authentication	MD5	
Dynamic Mesh	No	
Dynamic Mesh Timeout [sec]	600	
Security		
Traffic Intelligence		
TI - Bandwidth Protection		
Bandwidth Policy	Best Effort (no shaping)	
Assigned QoS Profile		
Estimated Bandwidth		
Low Priority		
TI - VPN Envelope Policy		
TOS Policy	Fixed Envelope TOS	
Envelope TOS Value	0	
QoS Policy	Use QoS Band from Host Ruleset	
QoS Band ID	1	
Advanced		
Key Time Limit [min]	10 mins	
Key Traffic Limit	No Limit	
Identification Type	Public Key	
Tunnel Probing [sec]	30 secs	
Tunnel Timeout [sec]	20 secs	
Packet Balancing	None	
WANOpt		
WANOpt Policy	NO-WANOpt	
GTI Settings		
Hide in Barracuda NG Earth	No	
Meshed	Yes	
Hub for this Group	HubVPN_CentralHub_2	
Service Placement	Classic circular	

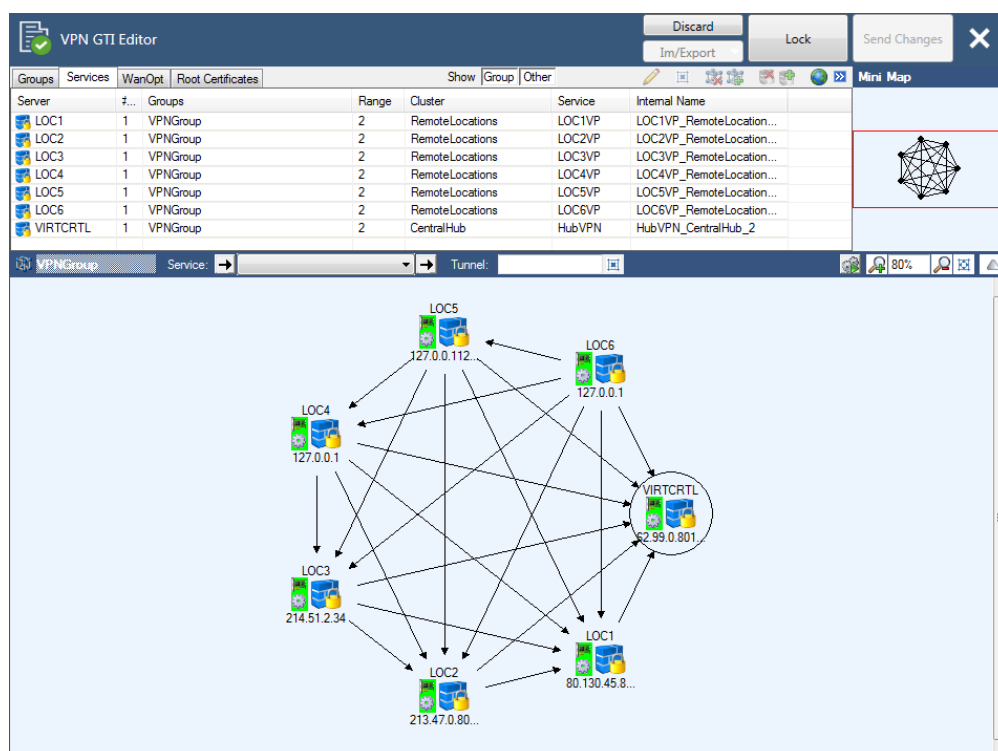
Add VPN Services to VPN Group

Add the VPN services to the VPN Group. If you are using the Range or Cluster GTI Editor, make sure to add only VPN services from that Range or Cluster to the group. Click on the HQ VPN service and select **Hub**. Since we selected the meshed option, the GTI Editor creates one VPN tunnel from each NextGen Firewall F-Series to all other firewalls in the VPN Group. If all the listening and transport IP addresses you configured for each VPN service were correct, all VPN tunnels will be up and running with a single transport using the first IP address in the Transport IP address list to establish the VPN tunnel. If the connection fails, the other IP addresses in the list are tried sequentially. In some cases, you have to adjust the configuration manually:


- For some VPN tunnels, it might be necessary to switch the active and passive tunnel partner. Delete the VPN tunnel and create a new VPN tunnel per drag and drop by starting at the VPN service that is now active.
- When two NextGen Firewall F-Series units are both using dynamic Internet connections, edit the passive side of the VPN tunnel and enter the DynDNS name of the active NextGen Firewall F-Series as an **Explicit Transport Listening** address. Be aware that the IP address of a FQDN is cached for the TTL of the domain by the VPN service. Depending on the timing of the DNS record update and TTL of the DynDNS record, it may not be possible to reconnect immediately.

You can [clear the DNS cache manually](#).

For each remaining tunnel status that is still red, log into the NextGen Firewall F-Series that initiates the VPN tunnel and check the VPN tab **Status** to gather more information. Adjust the settings for the tunnel as needed. Verify that the active NextGen Firewall F-Series can reach the transport listening IP address of the passive tunnel partner.



(Optional) Traffic Intelligence : Add Transports to the VPN Tunnels



Traffic Intelligence (TI) is the logical layer used to manage multiple parallel VPN tunnels (transports) in

one VPN tunnel configuration. Traffic Intelligence also handles loadbalancing, fail-over, and traffic routing for all transports of the VPN tunnel.

Without Traffic Intelligence, the VPN tunnel can use the bandwidth of only one of the Internet connections of the NextGen Firewall F-Series. So for firewalls using multiple WAN connections, add an additional transport for each connection to the VPN tunnels. Assign each transport a **TI Classification** and **TI ID**. When multiple IP addresses are used as **Transport Listening** IP addresses for one VPN service, the first IP address is always used to create the VPN tunnel. When that IP address is unavailable, the next IP addresses in the list are used until a VPN tunnel can be established. This behavior is undesired if you are using multiple, differently sized WAN connections. Instead, you should assign single, explicit **Transport Listening IP** addresses to each transport. You thereby know the available bandwidth for the VPN connection and can then assign matching traffic-shaping policies.

Through the **TI Learning Policy** settings you can determine which NextGen Firewall F-Series assigns the VPN transports for the connection (the **TI Master**). To avoid two TI masters overriding each other by sending traffic through different transports, the VPN hub must be configured to be the TI master while all remote firewalls are configured to be **TI Slaves**.

You control which transport is used for a specific connection, with the connection object of the access rule handling the VPN traffic. Configure weighted, session-based load balancing and fallback behavior for your transport. This gives you granular control over which transports and, by extension, which Internet connections are used for the VPN traffic. For example, you can configure your access rules so that VOIP uses an expensive, low-latency connection, whereas large transfers are delegated to the transport running on the cheaper Internet connection.

Transport Policies

Transport Selection Policy Explicit Transport Selection

T1 Learning Policy Master (propagate T1 settings to partner)

Explicit Transport Selection

Primary Transport Class Bulk

Primary Transport ID 0

Secondary Transport Class Bulk

Secondary Transport ID 0

Further Transport Selection First try Cheaper then try Expensive

☒ Allow Bulk Transports ☒ Allow Quality Transports ☒ Allow Fallback Transports

Simultaneous Transport Usage

Session Balancing None

Traffic Duplication No

TCP Transport Traffic Prioritization

When using BULK Transports High Priority

When using QUALITY Transports High Priority

Dynamic Mesh

☐ Allow Dynamic Mesh ☐ Trigger Dynamic Mesh

☒ Prevent Tunnel Timeout

Transport Policies

Transport Selection Policy

Explicit Transport Selection

TI Learning Policy

Slave (learn TI settings from partner)

Explicit Transport Selection

Primary Transport Class

Bulk

Primary Transport ID

0

Secondary Transport Class

Bulk

Secondary Transport ID

0

Further Transport Selection

First try Cheaper then try Expensive

☒ Allow Bulk Transports

☒ Allow Quality Transports

☒ Allow Fallback Transports

Simultaneous Transport Usage

Session Balancing

None

Traffic Duplication

No

TCP Transport Traffic Prioritization

When using BULK Transports

High Priority

When using QUALITY Transports

High Priority

Dynamic Mesh

☐ Allow Dynamic Mesh

☐ Trigger Dynamic Mesh

☒ Prevent Tunnel Timeout

For more information, see [Traffic Intelligence](#).

Section 5 - Firewall Services and Access Rules

Choosing the Right Firewall Service

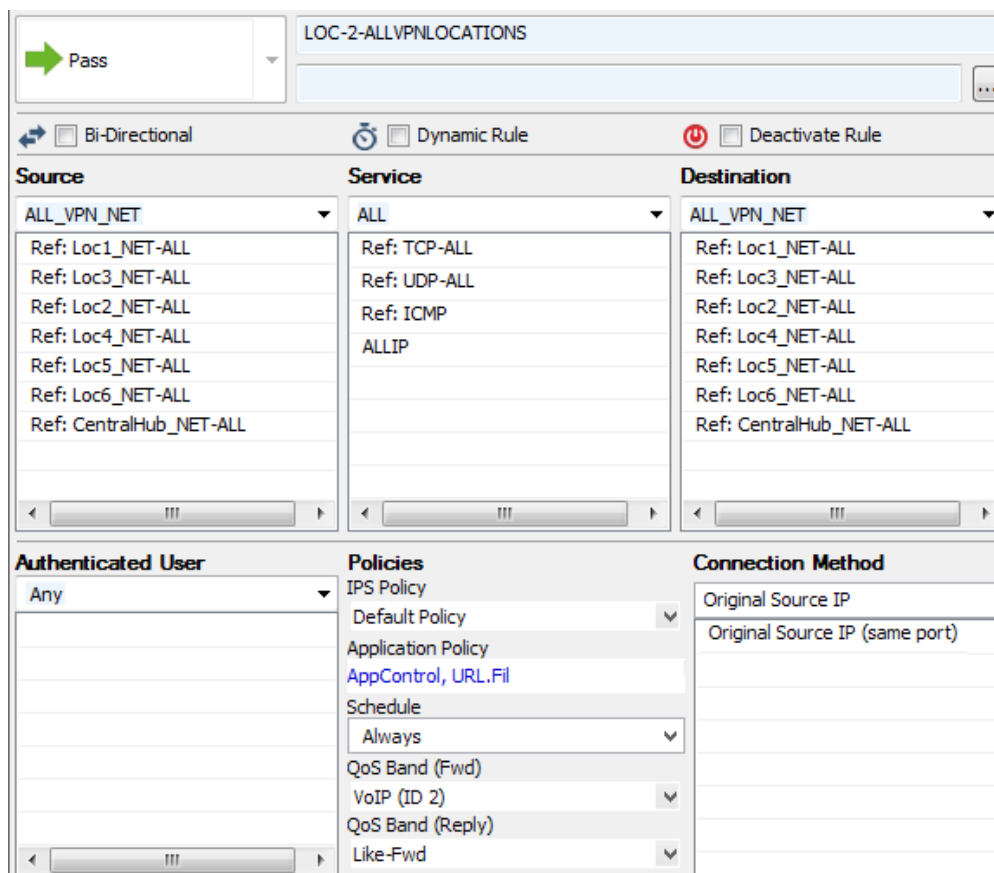
Access rules control which traffic is allowed in and out of a VPN tunnel. If you are using a meshed VPN network, you must take into account that traffic originating from every location can be sent through the VPN tunnel. You can reduce the configuration overhead by using the Distributed Firewall service for a cluster of NextGen Firewall F-Series units that share many access rules. Each location can still define specific rules in the **Local Rules** and **Special Rules**, but the admin can manage a common set of access rules in the **Global Rules**. If the NextGen Firewall F-Series units are not in the same cluster or do not share access rules, use a normal Forwarding Firewall service instead. For this example, we are using the Forwarding Firewall service for the central hub and a Distributed Firewall service for all remote locations.

For more information, see [Forwarding Firewall](#) and [Distributed Firewall](#).

Access Rule to Allow Traffic in and out of the VPN Tunnels

You need to create an access rule to allow traffic in and out of the VPN tunnels. This rule allows transparent access from all networks to all networks. Use this rule to validate that all the VPN networks are accessible and then substitute them with more specific access rules as necessary.

- **Action – PASS**
- **Source – ALL_VPN_NET** This is the network object we created containing all the networks in all locations.
- **Service – ALL** Create a service object containing all services you need to access through the tunnels.
- **Destination – ALL_VPN_NET** This is the network object we created containing all the networks in all locations.
- **Connection Method – Original Source IP** Replace with a custom connection object depending on what type of tunnel you want to create (see below).



The screenshot shows the configuration for an access rule named "LOC-2-ALLVPNLOCATIONS". The rule is set to "Pass" (indicated by a green arrow icon). It is a "Bi-Directional" rule, is not a "Dynamic Rule", and is not "Deactivated".

Source	Service	Destination
ALL_VPN_NET	ALL	ALL_VPN_NET
Ref: Loc1_NET-ALL	Ref: TCP-ALL	Ref: Loc1_NET-ALL
Ref: Loc3_NET-ALL	Ref: UDP-ALL	Ref: Loc3_NET-ALL
Ref: Loc2_NET-ALL	Ref: ICMP	Ref: Loc2_NET-ALL
Ref: Loc4_NET-ALL	ALLIP	Ref: Loc4_NET-ALL
Ref: Loc5_NET-ALL		Ref: Loc5_NET-ALL
Ref: Loc6_NET-ALL		Ref: Loc6_NET-ALL
Ref: CentralHub_NET-ALL		Ref: CentralHub_NET-ALL

Authenticated User	Policies	Connection Method
Any	IPS Policy	Original Source IP
	Default Policy	Original Source IP (same port)
	Application Policy	
	AppControl, URL.Fil	
	Schedule	
	Always	
	QoS Band (Fwd)	
	VoIP (ID 2)	
	QoS Band (Reply)	
	Like-Fwd	

Using Connection Objects to Hide Networks behind VPN Tunnels

Depending on the connection object used for the access rules allowing traffic in and out of the VPN

tunnels, you can:

- Hide all/some remote clients behind one IP address.
- Allow completely transparent network access between the locations.
- Select the transport for TI with session-based loadbalancing and failover.

For more information, see [Examples for TINA VPN Tunnels](#).

Distributed Firewall Service: Create Access Rules for the Remote Locations

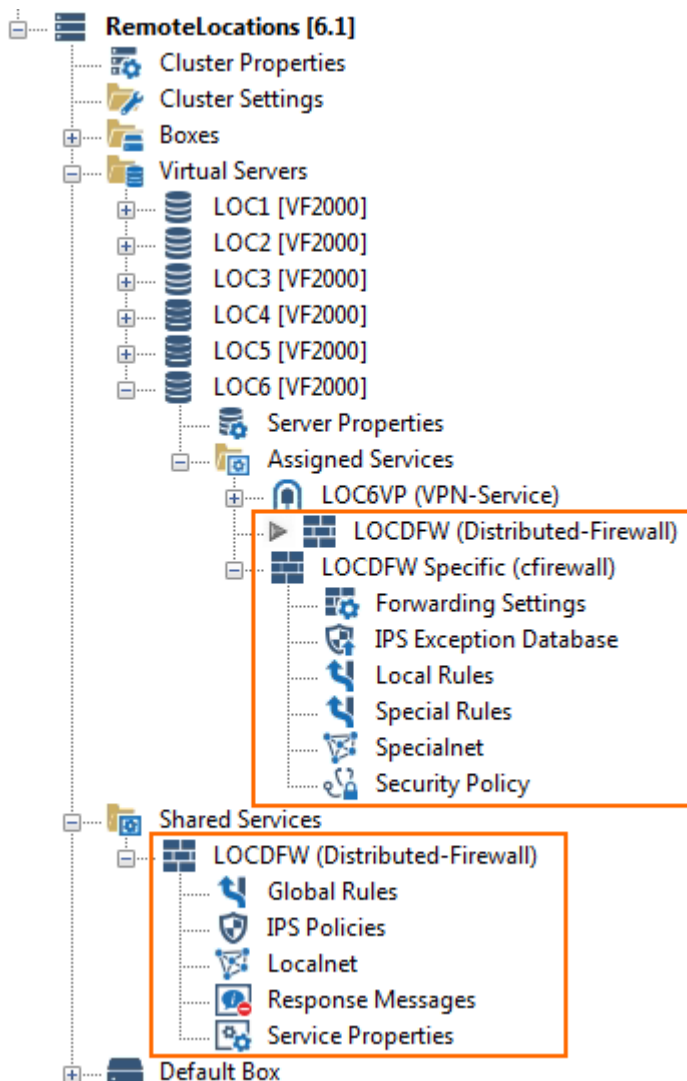
The Distributed Firewall service splits the firewall ruleset into a global ruleset, which is valid for all firewalls using the shared service, and a local and special ruleset. You must create **Cascade** access rules in the global ruleset for these access rules to be evaluated.

In the Global ruleset, create the following rules:

- Add a Cascade to the local ruleset at the beginning of the global ruleset.
- Add a rule allowing traffic in and out of the VPN tunnels (see above). Use the ALL_NET global network object.

In the Local ruleset, create the following rules:

- Add special rules, e.g., VPN tunnel access rules with custom connection objects, DHCP to VPN service redirect rule for dynamic Internet connections, ...
- Add a Cascade back rule to each local ruleset before the Block All rule.



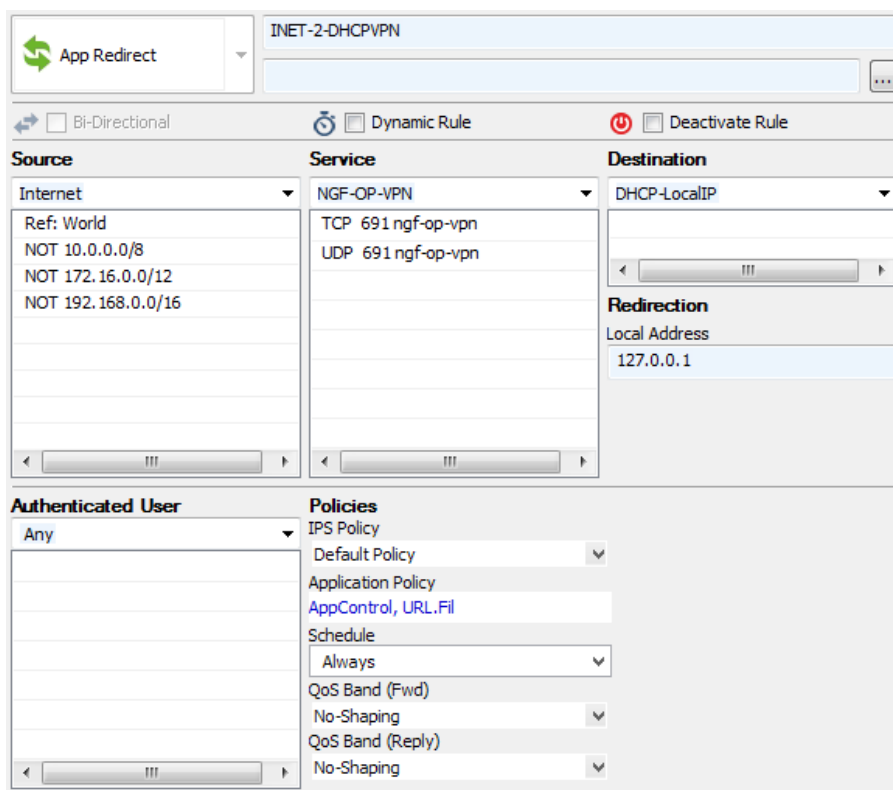
Global Firewall Rules									
Main Rules									
Action	Name	Features	Service	Source	Destination	User	Sched...	QoS	IPS Policy
0 Cascade local	CASC-2-LOCAL-Ruleset		ALL ALLIP, ECHO, TCP *	Any 0.0.0.0/0	Any 0.0.0.0/0	Any	Always	N.A.	N.A.
Common Rules (5)									
3 Pass No Src NAT [Client]	LAN-2-LAN		ALL ALLIP, ECHO, TCP *	Trusted LAN 0.0.0.0	Trusted LAN 0.0.0.0	Any	Always	No-Shaping	0 Default Policy
4 Pass No Src NAT [Client]	LOC-2-ALLVPNLOCATIONS		ALL ALLIP, ECHO, TCP *	ALL_VPN_NET 10.0.10.0/25, 10.0.15.0/2...	ALL_VPN_NET 10.0.10.0/25, 10.0.1...	Any	Always	No-Shaping	0 Default Policy
5 Pass Dynamic Src NAT [Proxydyn]	LAN-2-INTERNET		ALL ALLIP, ECHO, TCP *	ALL_VPN_NET 10.0.10.0/25, 10.0.15.0/2...	Internet 0.0.0.0/0, NOT 10.0....	Any	Always	No-Shaping	0 Default Policy
6 Cascade special	CASC-2-SPECIAL-Ruleset		ALL ALLIP, ECHO, TCP *	Any 0.0.0.0/0	Any 0.0.0.0/0	Any	Always	N.A.	N.A.
7 Block	BLOCKALL		ALL ALLIP, ECHO, TCP *	World 0.0.0.0/0	World 0.0.0.0/0	Any	Always	N.A.	N.A.

Special Considerations for NextGen Firewall F-Series Units with a Dynamic Internet Connection

The NextGen Firewall F-Series units using a dynamic Internet connection must redirect all incoming VPN traffic to the VPN service running on the 127.0.0.1 IP address.

Create an App Redirect access rule to redirect incoming VPN traffic to your VPN service listening on 127.0.0.1

- **Action** – App Redirect
- **Source** – Internet
- **Service** – **NG-OP-VPN**. If you are using ESP as the transport mode, you must also add ESP to the service.
- **Destination** – Select the dynamic network object that matches your Internet connection. For example, **DHCP-LocalIP** for a DHCP Internet connection.
- **Redirection** – Enter 127.0.0.1.



The screenshot shows the configuration for an 'App Redirect' rule. The rule name is 'INET-2-DHCPVPN'. The 'Source' is set to 'Internet' (Ref: World, NOT 10.0.0.0/8, NOT 172.16.0.0/12, NOT 192.168.0.0/16). The 'Service' is 'NGF-OP-VPN' (TCP 691 ngf-op-vpn, UDP 691 ngf-op-vpn). The 'Destination' is 'DHCP-LocalIP'. The 'Redirection' section shows 'Local Address' as '127.0.0.1'. The 'Authenticated User' is 'Any'. The 'Policies' section includes 'IPS Policy' (Default Policy), 'Application Policy' (AppControl, URL.Fil), 'Schedule' (Always), 'QoS Band (Fwd)' (No-Shaping), and 'QoS Band (Reply)' (No-Shaping).

Section 6 - Additional Topics / Optimizations

Depending on your network and requirements, you can also use the following features to tailor your network to your needs:

Dynamic Mesh

A Dynamic Mesh VPN network allows you to use the advantages of a fully meshed network without having to provide the resources needed for the large number of static VPN tunnels on every unit. All remote units are connected by a static TINA VPN tunnel to a central NextGen Firewall F-Series acting as the VPN hub. When relay traffic from a remote NextGen Firewall F-Series to another remote

NextGen Firewall F-Series is detected by the VPN hub, a dynamic VPN tunnel is initiated between the two remote firewalls. As soon as the dynamic VPN tunnel is up, traffic is transparently redirected through the VPN tunnel that now directly connects both locations. The dynamic tunnel is completely transparent to the user and offers better latency than relaying the traffic through the VPN hub. Dynamic tunnels are triggered by the dynamic-mesh-enabled connection object of the VPN hub. Configure the VPN hub as the TI master, and the remote units as TI slaves. The TI slaves will automatically learn the Dynamic Mesh and TI settings from the master. Traffic that does not match an access rule with a dynamic-mesh-enabled connection object on the TI master continues to be sent through the VPN hub.

To use a dynamic mesh instead of a fully meshed network, create Site-to-Site tunnels only between the remote locations and the VPN hub, and use a dynamic-mesh-enabled connection object on the VPN hub to trigger dynamic, on-demand tunnels between the remote locations.

For more information, see [Dynamic Mesh VPN Networks](#).

WAN Optimization

WAN Optimization reduces the amount of traffic sent through the tunnel. You can attain very high deduplication rates depending on the type of traffic going through the tunnel and the amount of available CPU resources. However, WAN Optimization is less effective if you send a lot of UDP or encrypted TCP traffic. Also note that if you rely on SSL Interception, Virus Scanning, or ATP, these features do not work in combination with WAN Optimization.

For more information, see [WAN Optimization](#).

VPN Compression

If you cannot use WAN Optimization, you can alternatively enable VPN compression to save bandwidth with minimal configuration overhead. VPN compression is not as effective as WAN Optimization, but it can be used in combination with Application Control 2.0.

For more information, see [TINA Tunnel Settings](#).

Traffic Shaping (QoS)

Applying traffic shaping policies to VPN traffic can be configured in two ways, each with its own set of

advantages and limitations.

Shape on VPN Transports

Applying shaping policies directly to the VPN transports allows you to shape individual transports. A limitation of this approach is defining optimal inbound and outbound bandwidth settings for systems using multiple transports on one ISP connection is not possible. If the value is set too low, the transport cannot use all the potentially available bandwidth on the network interface. Setting the value too high may cause the available bandwidth of the network interface to be exceeded, causing the traffic shaping engine to drop random VPN packets. This issue does not occur when using consolidated traffic shaping.

Consolidated Traffic Shaping

Consolidated traffic shaping shapes the traffic inside a VPN tunnel with the settings of the network interface used to send the VPN traffic. This lets you define policies to prioritize important traffic if it is sent directly through an interface or is encapsulated into a VPN transport.

Since traffic shaping is applied before either VPN compression or WAN Optimization, using consolidated shaping may result in unused bandwidth on the network interface. This issue does not occur with transport-based traffic shaping because the shaping engine for the physical network interface uses the compressed VPN tunnel packets.

For more information, see [Traffic Shaping](#) and [How to Apply Traffic Shaping to a VPN Tunnel](#).

Figures

1. IG_00.png
2. IG_00a.png
3. all_fws_empty.png
4. allfws_clusters_ranges.png
5. IG_10_GlobalFWObjects.png
6. IG_11_GlobalFWObjects.png
7. IG_09_RepositoryLinkedException.png
8. IG_08_Repository.png
9. all_fws_services.png
10. loc1_blowup_net_objects.png
11. loc6_dhcp_gti.png
12. IG_03.png
13. IG_02.png
14. loc1_2_ti_settings.png
15. IG_07_TI_Master.png
16. IG_06_TI_Slave.png
17. IG_04.png
18. IG_12_DFW_Structure.png
19. IG_13_DFW_Global_Rules.png
20. IG_05.png

© Barracuda Networks Inc., 2024 The information contained within this document is confidential and proprietary to Barracuda Networks Inc. No portion of this document may be copied, distributed, publicized or used for other than internal documentary purposes without the written consent of an official representative of Barracuda Networks Inc. All specifications are subject to change without notice. Barracuda Networks Inc. assumes no responsibility for any inaccuracies in this document. Barracuda Networks Inc. reserves the right to change, modify, transfer, or otherwise revise this publication without notice.