# Dynamic Mesh VPN Networks

https://campus.barracuda.com/doc/53248939/

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 the central firewall acting as the VPN hub. VPN tunnels between IPv6 endpoints are not supported. When relay traffic from a remote location to another remote location is detected by the VPN hub, a dynamic VPN tunnel is initiated directly connecting 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 prevent services such as OSPF or BGP from keeping dynamic tunnels open forever, you can disable resetting the idle timeout of the dynamic tunnel in the connection object of the matching access rule.

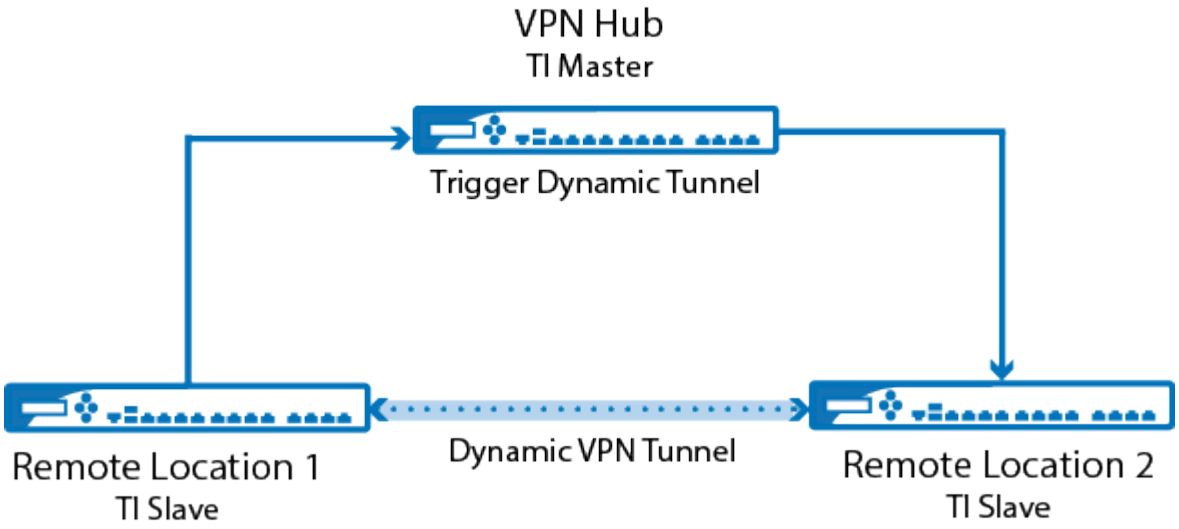


## Video

Watch the following video to see a Dynamic Mesh VPN in action


## Initiating a Dynamic Tunnel

A dynamic tunnel is created when the following requirements are met:

- All firewalls must use IPv4 transport source and listening IP addresses.
- Both firewalls must be connected to the same VPN hub via TINA VPN tunnels.
- The VPN hub must act as a relay. For example, traffic must pass through the VPN hub to the target NextGen Firewall F-Series.
- The VPN hub must be configured as the TI master.
- The remote firewalls must be configured as TI slaves.
- The source NextGen Firewall F-Series must be able to reach the public IP address of the target NextGen Firewall F-Series. If multiple VPN listening IP addresses are present, the first IP address from the list is chosen.
- Dynamic Mesh must be enabled on each NextGen Firewall F-Series and the VPN hub in the **VPN Settings**
- The VPN hub acting as the TI master must have **Allow Dynamic Mesh** and **Trigger Dynamic Mesh** enabled in the connection object.
- The tunnel is terminated if no traffic is sent through the tunnel for the configured timeout. (Min:10 sec. Default 600 sec.)

## Dynamic Tunnel Settings

Ideally, both VPN tunnels connecting to the hub use the same encryption and transport settings. If these settings differ, the dynamic tunnel uses the following preferences:

- **Transport** – If the **Transport** settings differ, the dynamic tunnel chooses the transport protocol according to the following preferences:
  1. ESP
  2. UDP
  3. TCP
- **Compression** – Compression is enabled for the dynamic tunnel if at least one of the static

tunnels also uses compression.

- **Encryption** – If the **Encryption** settings differ, the dynamic tunnel chooses the cipher according to the following preferences:
  1. AES
  2. BLOW
  3. CAST
  4. 3DES
  5. DES
  6. NONE
- **Authentication** – If the **Authentication** settings differ, the dynamic tunnel chooses the hash according to the following preferences:
  1. GCM
  2. SHA512
  3. SHA256
  4. MD160
  5. SHA
  6. MD5
  7. NONE

### Traffic Intelligence with Dynamic Mesh Tunnels

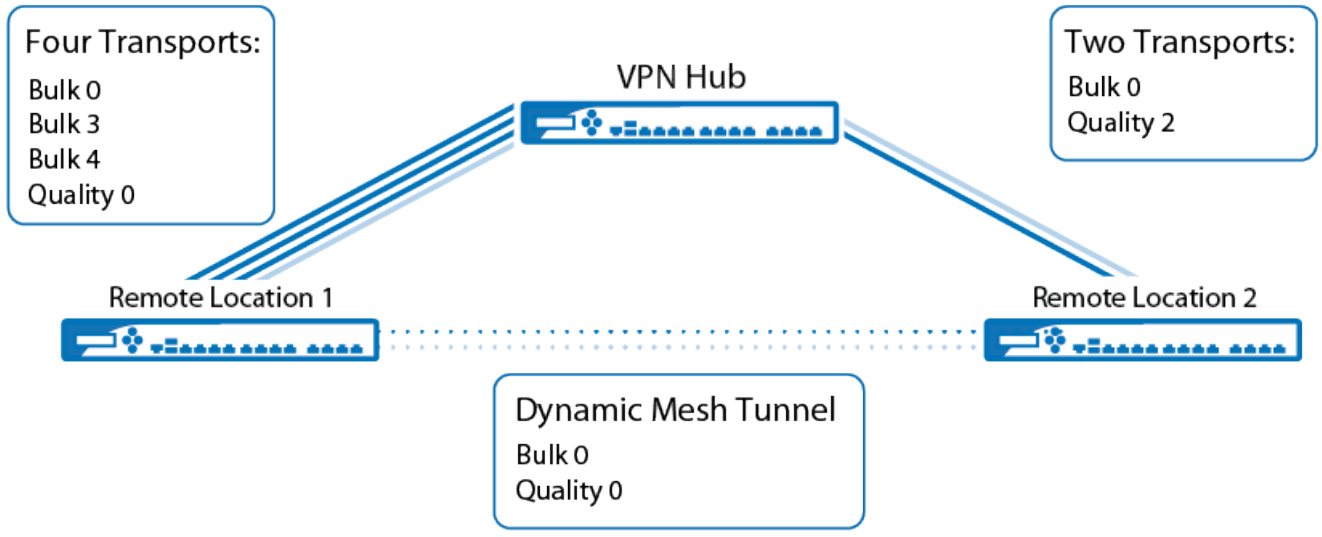


When a dynamic tunnel is created between two F-Series Firewalls both using multiple transports, the dynamic tunnel will create a transport with the TI ID of 0 for **Bulk** and **Quality** TI classes used in at least one of the static VPN tunnels. This means that for two remote VPN services using multiple transports in the TI class bulk, the dynamic tunnel will be created with a single Bulk0 transport. The source networks from the static tunnels are assigned to the transports of the dynamic tunnel according to their TI class. For example, if a network was previously routed through the bulk3

transport, it will be assigned to the bulk0 transport of the dynamic tunnel. The VPN hub must act as TI Master, and the remote units as TI slaves. The remote firewalls will learn the dynamic mesh settings from the TI master. When two TI slaves communicate with each other, the transport is chosen by the **TI Transport Selection** configured for the connection object of the NextGen Firewall F-Series initiating the connection. Make sure the **Fallback** policy allows the use of the TI ID 0 of each transport. It is recommended to use identical firewall connection objects for all remote firewalls.

## Limitations

- Dynamic Mesh cannot be used in combination with WAN Optimization.
- Traffic Shaping must be applied to the VPN interface and not directly to the transport.
- Dynamic Mesh cannot be used for F-Series Firewalls that are behind a NATed connection which hinders the VPN hub from finding out the public IP address of the remote unit.
- VPN tunnel start/stop scripts are not executed on the remote F-Series Firewalls
- Dynamic tunnels are not synced to the HA partner.

## Dynamic Mesh Configuration via GTI Editor for Managed F-Series Firewalls

The GTI Editor simplifies configuring a large Dynamic Mesh VPN network for firewalls managed by a Control Center.

For more information, see How to Configure a Dynamic Mesh VPN with the GTI Editor.

## Dynamic Mesh Configuration on Stand-alone F-Series Firewalls

Dynamic Mesh can be configured for VPN networks with three or more stand-alone firewalls, with the central NextGen Firewall F-Series acting as a VPN relay and hub.

For more information, see How to Configure Dynamic Mesh VPN.

## Figures

1. ti_dyn_mesh_01.png
2. ti_dyn_mesh_02.png