

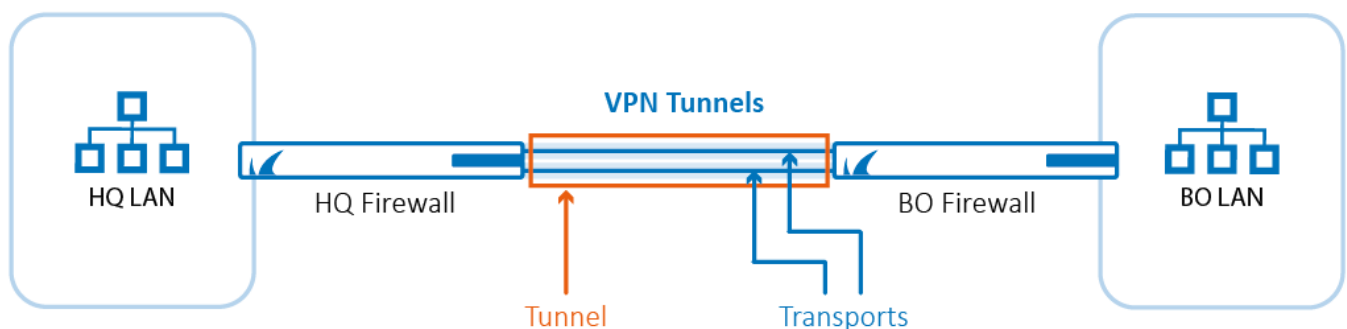
Traffic Intelligence

<https://campus.barracuda.com/doc/79462906/>

Traffic Intelligence (TI) provides multiple VPN transports with each transport capable of using a different WAN connection, thereby expanding on the concept of a traditional VPN tunnel with only one VPN transport to one logical VPN tunnel. TI also provides redundant, reliable, and failsafe network connections: the VPN tunnel is up and can transmit traffic as long as at least one transport is operational. Admins can retain full control over how each transport is used, or they can configure the advanced balancing and bandwidth management features to optimally use the available bandwidth. Note that since TI requires the TINA VPN protocol, both the local and remote gateway must be Barracuda CloudGen Firewalls. Traffic Intelligence combines a multi-transport VPN tunnel with the following advanced VPN routing, balancing, and shaping features:

- VPN Transports
- Dynamic Bandwidth and Latency Detection
- Performance-Based Transport Selection
- Adaptive Bandwidth Protection
- Adaptive and Static Session Balancing
- Traffic Duplication

VPN Transports




When connecting two sites, a single transport tunnel can use only one WAN connection for each site. Therefore, to use multiple WAN connections, multiple parallel VPN tunnels would have to be created, resulting in difficulties when routing traffic over these parallel tunnels. However, by using multiple transports, only one VPN tunnel and the routes for one tunnel are needed. For each WAN connection, a VPN transport is added to the VPN tunnel. The connection object of the access rule that matches traffic determines which transport is used. Transports can use a mix of IPv4 and IPv6 WAN connections, MPLS lines, and fallback WWAN connections. The transport protocol used can be set individually for each VPN transport, depending on the type of traffic and WAN connection: UDP, TCP, ESP, or Routing. Transports are split into three classes, with each class containing up to eight IDs for a maximum total of 24 transports per VPN tunnel.

VPN Transport Classes

The three VPN transport classes are classified according to their "cost":

- **Bulk** - For cheap and potentially unreliable connections. Bulk transports are recommended for xDSL or cable WAN connections.
- **Quality** - For a more reliable line, such as a business-quality Internet line or MPLS links.
- **Fallback** - For the most expensive lines. Fallback transports are recommended for dial-in lines or WWAN connections.

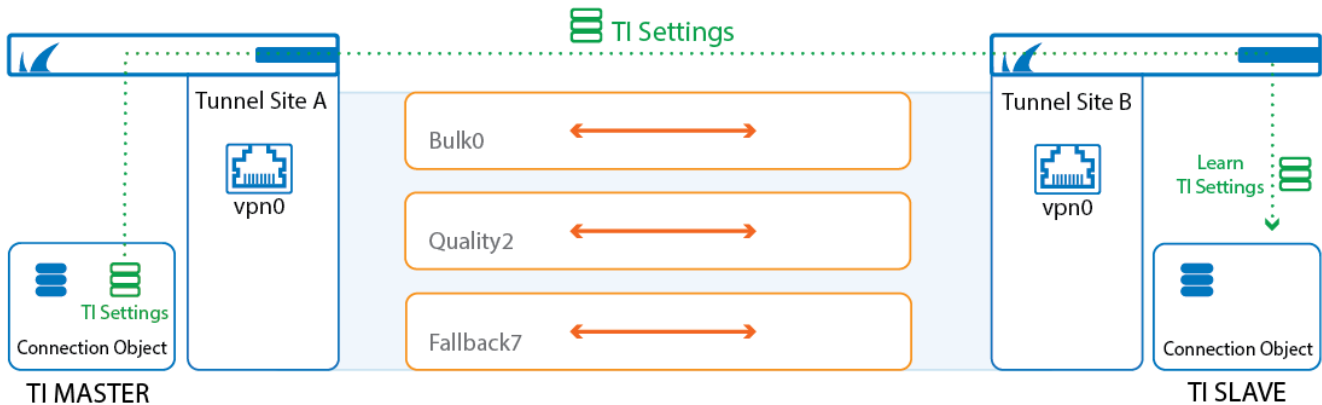


VPN Transport Class IDs

Each VPN transport class is made up of eight class IDs (**0 - 7**), which define the VPN transport cost in more detail. The class IDs provide you with more configuration options for creating VPN transports in a single VPN tunnel. A higher metrics indicates a more expensive transport.

Bulk								Quality								Fallback								TI-Class
0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	TI-ID

Create Multi-Transport TINA VPN Tunnels

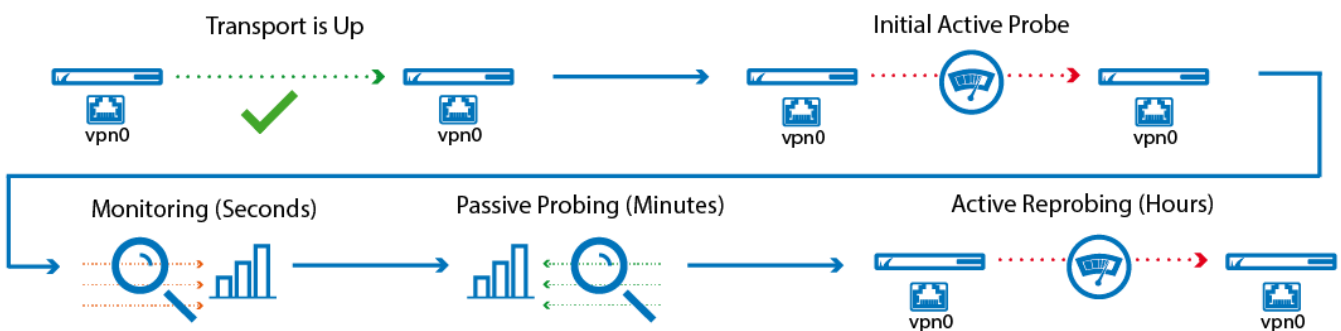


Multi-transport VPN tunnels can be configured either manually for each TINA site-to-site VPN tunnel, or via the GTI Editor if both firewalls are managed by the same Firewall Control Center. The TI settings of the custom connection object used in the matching access rule determines which transport is used. In addition to transport balancing, failover, and advanced bandwidth management, features can be enabled to fully utilize all available WAN connections.

Traffic is routed through a VPN transport by the TI settings of the connection object in the matching access rule. The TI settings allow for simple, one-transport routing, as well as complex, adaptive balancing between different transports. To ensure that the same TI settings are always used by both tunnel endpoints, one firewall is the TI master, the other the TI slave. The TI master propagates the TI settings, overwriting the TI settings on the TI slave.

For more information, see [How to Add a VPN Transport to a TINA VPN Tunnel with Explicit Transport Selection](#) and [How to Configure Traffic Intelligence Using the VPN GTI Editor](#).

Dynamic Bandwidth and Latency Detection



For UDP transports, the firewall can determine the actual bandwidth available for a VPN transport through monitoring, active probing, and passive probing. Dynamic bandwidth and latency detection is available for IPv4 connections only. To have a valid starting point, the initial bandwidth is set in the VPN transport configuration. The goal for the link-quality probing is to find the settings that offer the best possible combination of latency and bandwidth with the fewest dropped packages. To determine

the effective bandwidth, the firewall compares, among other things, the number packets sent and received at either end of the VPN. This also yields the number of dropped packets and, at the same time, the latency (round-trip time) of each transport.


1. Initial Active Probing

- The VPN transport is established.
- After a couple of seconds, the initial active probe is started. The expected bandwidth entered by the admin is used as the starting point.
- The bandwidth, latency, and drop rate are applied to the transport.

2. Monitoring

Granularity for monitoring is measured in seconds.


- Latency, drops, and bandwidth are continuously monitored as traffic passes through the transport.
- Monitoring is used to detect lower bandwidth.



3. Passive Probing

Executed every couple of minutes.

- Passive probing to detect increases in the available bandwidth. Traffic already using the transport is not influenced by probing.



4. Active Reprobe

Executed every couple of hours.

- A repeat of the initial active probe.

Dynamic Bandwidth and Latency Detection does not have to use the full probing and monitoring solution to determine the link quality. If the quality of the link is very stable, it may make sense to reduce the probing and monitoring, or to disable it altogether and use static values for the bandwidth instead:

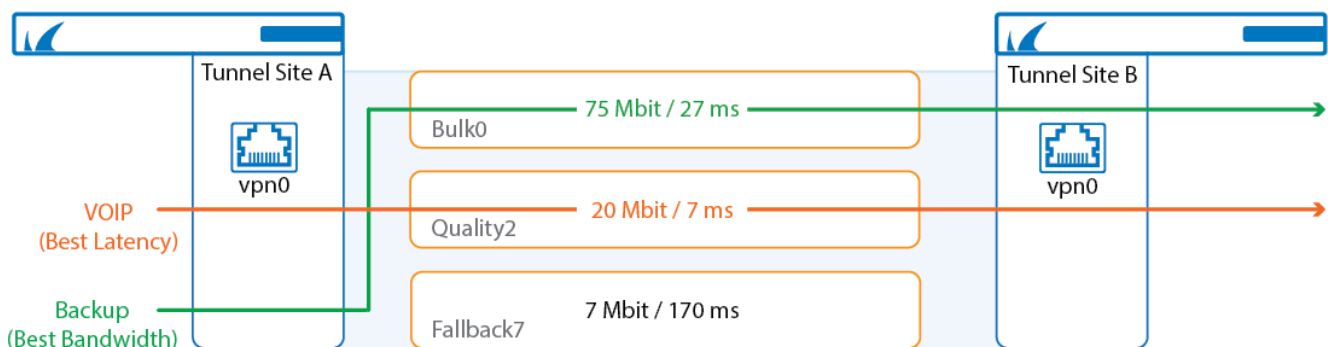
- **Active Probing and Passive Monitoring** – All probing and monitoring features are used to determine the link quality metrics.
- **Active Probing Only** – The initial active probe and the hourly active reprobe are used to determine the link quality metrics.
- **No Probing - use Estimated Bandwidth** – Probing is disabled. Features using Dynamic

Bandwidth and Latency Detection use the estimated bandwidth entered by the admin in the VPN tunnel configuration.

Enabling or disabling the Dynamic Bandwidth and Latency Detection requires a manual termination of the VPN transport to take effect. This is not required for changing the Dynamic Bandwidth and Latency Detection modes.

Dynamic Bandwidth and Latency Detection is required to be able to use Adaptive Bandwidth Protection, and Adaptive Session Balancing. It is not possible to use these features in combination with TCP, ESP, and hybrid transport protocols. Dynamic Mesh VPN is also not supported.

Performance-Based Transport Selection



Performance-Based Transport Selection selects the optimal transport based on the policy selected in the TI settings of the custom connection object. Only UDP transports with Dynamic Bandwidth and Latency Detection enabled are included in the Performance-Based Transport Selection policy. The transport selections are made from the point of view of the TI master. The following policies are available:

- **Optimize for Latency** – Traffic is sent through the VPN transport with the lowest latency. If the latency changes, the selected transport is also updated.
- **Optimize for Inbound Bandwidth** – Traffic is sent through the VPN transport with the highest available downstream bandwidth for the QoS class from the TI master's point of view. No-delay traffic uses the total bandwidth as the criteria. Standard traffic uses the total bandwidth minus the no-delay traffic to make the decision of which transport to use.
- **Optimize for Outbound Bandwidth** – Traffic is sent through the VPN transport with the highest available upstream bandwidth for the QoS class from the TI master's point of view. No-delay traffic uses the total bandwidth as the criteria. Standard traffic uses the total bandwidth minus the no-delay traffic to make the decision of which transport to use.
- **Optimize for Combined Bandwidth** – Traffic is sent through the VPN transport with the highest bandwidth calculated by adding the upstream and downstream bandwidths from the TI

master's point of view. If other traffic is also using this transport, this might not correlate with the highest available bandwidth. Again, the same logic applies for no-delay and standard traffic. No-delay traffic uses the total combined bandwidth as the criteria. Standard traffic uses the total bandwidth minus the no-delay traffic on the transport to make the decision.

For more information, see [How to Configure Performance-Based Transport Selection for VPN Tunnels with Traffic Intelligence](#).


Traffic Shaping for VPN Transports

Traffic shaping can be applied to VPN traffic in different ways. You can shape the output interface or the VPN transport directly. Determining the best method is not always easy since assigning a static bandwidth to a transport is difficult if there is other traffic on the same WAN interface. For UDP transports, the firewall can also use a series of dynamic link quality checks to determine the optimal bandwidth at any given moment. This allows the firewall to react to changes to latency, bandwidth, or packet drops.

Simplified Traffic Shaping Tree for Adaptive Bandwidth Protection

To use Traffic Intelligence with the adaptive bandwidth features, a simplified traffic shaping tree is used. Traffic can be classified either as nodelay or standard. 30% of the bandwidth is always reserved for NoDelay traffic, and if there is no standard traffic can take up to 100% of the available bandwidth. Standard traffic can take up to 70% of the available bandwidth leaving the 30% reserved bandwidth for nodelay traffic free. If both nodelay and standard traffic are present, nodelay traffic can take up to 90% of the available bandwidth, leaving 10% for standard traffic on a single transport.

The traffic shaping tree used by Adaptive Bandwidth Protection is not visible in the QoS profile tab of the Traffic Shaping configuration and cannot be changed or replaced by the user. If you are using the default QoS profiles and bands use the **VoIP** and **LowPrio**. If you have customized the existing default QoS bands it is recommended to create two additional bands to classify standard and NoDelay traffic. For NoDelay traffic, create a QoS band using priority class **NoDelay** on the root interface. For the QoS band for standard traffic, create a QoS band using a virtual interface. Although any virtual interface will work, it is recommended to create a dedicated STD virtual interface to be able to read the configuration better. All settings of the virtual interfaces are handled by the firewall when used for the SD-WAN Traffic Intelligence features.



Adaptive Bandwidth Protection

Adaptive Bandwidth Protection ensures that traffic in the NoDelay (VoIP) QoS band is always prioritized over standard traffic. The firewall uses the link quality metrics gathered by Dynamic Bandwidth and Latency Detection to adjust traffic shaping to always fully utilize the available bandwidth. Passive monitoring allows the firewall to detect decreases in bandwidth; active probes increase the traffic sent through the link to determine if the bandwidth of the transport can be increased. The VPN monitoring graph displays these active probes as short spikes. Large jumps in quality might require multiple probes before you can determine the correct bandwidth for transport. It is recommended to combine Adaptive balancing on the VPN transport with consolidated shaping to shape the VPN traffic in a two step process:

- **Adaptive Shaping on the VPN Transport** - Shapes on the transport with a focus on site-to-site traffic in one VPN tunnel. For example: backup and voice traffic on the same VPN transport.
- **Consolidated Shaping** - Shapes the VPN traffic as a whole. Consolidated shaping is best used to control simultaneous traffic from many sites. This protects standard traffic from one VPN crowding out nodelay traffic on another VPN tunnel.

For more information, see [How to Configure Adaptive Bandwidth Protection for VPN Tunnels with Traffic Intelligence](#).

Static Shaping

For VPN transports using TCP, ESP, or hybrid transport protocols, only static shaping is available. You can shape on the transport of the VPN output interface using static values estimating the available bandwidth. Setting this value too high renders shaping useless, whereas setting it too low makes the bandwidth utilization less efficient.

Transport Balancing

Transport balancing is used to combine load balancing sessions of packages across multiple VPN transports, thereby effectively increasing the available bandwidth for this type of traffic. Otherwise, the transport is statically assigned in the connection object, which is distributed across two or more transports. Load balancing is completely transparent to the user. Two types of load balancing are available:

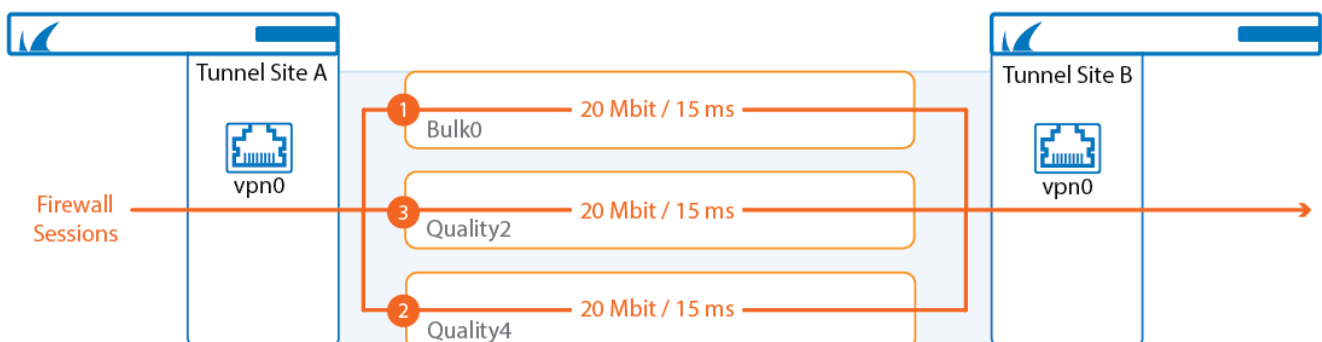
- Session Balancing
- Packet Balancing

Session Balancing

Session balancing distributes VPN traffic over multiple transports. It can be configured in two modes:

- **Static Session Balancing**
- **Adaptive Session Balancing**

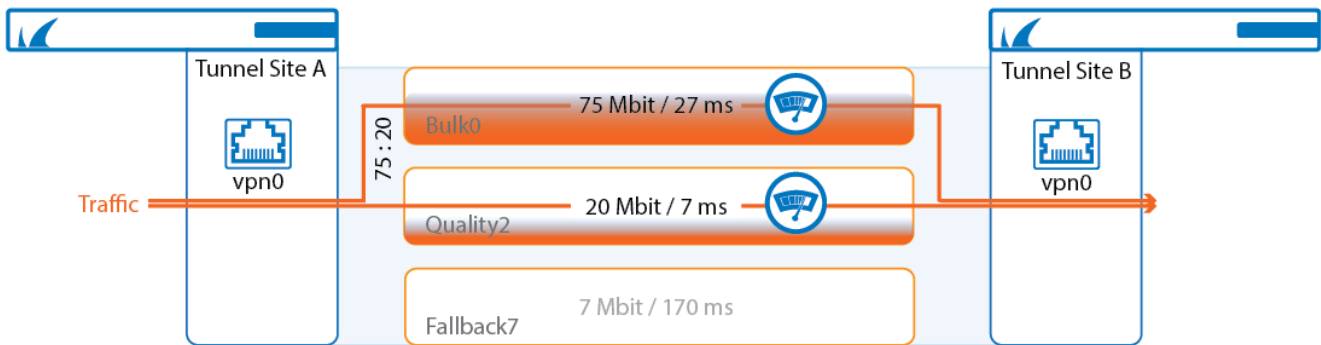
Static Session Balancing distributes all firewall sessions via round robin over the selected transports without regard to the available bandwidth on each individual transport. Session balancing must be enabled in the TI settings of the connection object in the matching access rule. When used without adaptive session balancing, it is recommended to use transports of roughly the same bandwidth. Static session balancing is supported for all VPN transport protocols (UDP, TCP, hybrid, and routing). Static session balancing can be configured to balance over just the primary and secondary transports or multiple transports in the same TI class based on the TI ID range defined in the connection object.



Adaptive Session Balancing uses link-quality metrics collected by Dynamic Bandwidth and Latency Detection for both the initial balancing and to rebalance sessions with a lifetime over 5 seconds. Adaptive balancing can only be configured between the primary and secondary transport. Only


transports using UDP as the transport protocol can be used. When selecting the transport, the firewall also takes asymmetric links into account, selecting the transport that offers the best upstream or downstream performance based on the selected balancing policy. Sessions shorter than 5 seconds stay on the initial transport and are not rebalanced. Rebalancing happens continuously, to always select the optimal transport.

When combined with Adaptive Bandwidth Detection, transport selection also takes the QoS band and the available bandwidth for the QoS band into account. NoDelay traffic uses the detected bandwidth of the transport to calculate which transport is chosen. Standard traffic subtracts the NoDelay traffic from the detected bandwidth before deciding on the transport. This way standard traffic is not assigned a transport that is already filled up with NoDelay traffic.



For more information, see [How to Configure Session Balancing for VPN Tunnels with Traffic Intelligence](#).

Packet Balancing




Packet-based balancing requires transports with the same latency and bandwidth, for example multiple identical WAN links from the same ISP. The VPN traffic is balanced with a round robin balancing policy on a per-packet basis over multiple VPN transports. Packet-based balancing is enabled in the TINA VPN tunnel configuration.

In most cases, it is recommended to use (adaptive) session-based balancing because it offers more flexibility and is more tolerant of differing link qualities.

For more information, see [How to Configure Packet-Based Balancing for VPN Tunnels with Traffic Intelligence](#).

Traffic Duplication

Traffic duplication copies packets and simultaneously sends them through the selected primary and secondary transports. Both traffic streams are combined again at the other end of the VPN tunnel. Use traffic duplication for applications requiring instant failover without a single dropped packet in case a VPN transport goes down. Since traffic is duplicated, both transports must have the same bandwidth and latency.



For more information, see [How to Configure Traffic Duplication for VPN Tunnels with Traffic Intelligence](#).

Figures

1. TI_01.png
2. ti_classification.png
3. TI_02.png
4. ti_learning.png
5. ti_probing.png
6. probing_monitoring.png
7. probing_passthrough.png
8. ti_performance_based_transport_selection.png
9. sdwan_traffic_shaping_tree.png
10. ti_session_balancing.png
11. ti_adaptive_balancing.png
12. ti_packet_balancing.png
13. ti_traffic_replication.png

© Barracuda Networks Inc., 2019 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.