



SpeedFusion and Best Practices

How to Get the Most out of SpeedFusion Technologies

Introduction to SpeedFusion

The purpose of this whitepaper is to explain how SpeedFusion works and to detail best practices for SpeedFusion SD-WAN deployment so that partners and end users can get the most out of SpeedFusion-enabled Peplink and Pepwave devices.

The intended audience is current or potential customers and value added resellers who are deploying and integrating SpeedFusion-enabled devices in their networks.

Technology Overview

SpeedFusion is a proprietary, patented VPN bonding technology. Since this is a VPN technology and VPN is a point-to-point protocol, two or more SpeedFusion-enabled devices are required, just as you'd expect.



One of SpeedFusion's most powerful features is that it can use multiple WAN links to create a single logical VPN tunnel between devices. This allows SpeedFusion bonding to provide two key functions: bandwidth aggregation and VPN reliability.

All WAN links connected to a SpeedFusion-enabled device can simultaneously work together as a single logical VPN connection. If a link fails, SpeedFusion can detect this failure and seamlessly redirect traffic, at a packet level, across other available links. This WAN link failure detection at the packet level allows SpeedFusion to provide VPNs with highly reliable and resilient site-to-site connectivity.

Probably the World's Easiest VPN

How to set up a site to site VPN using InControl Central Cloud management on Peplink Routers:

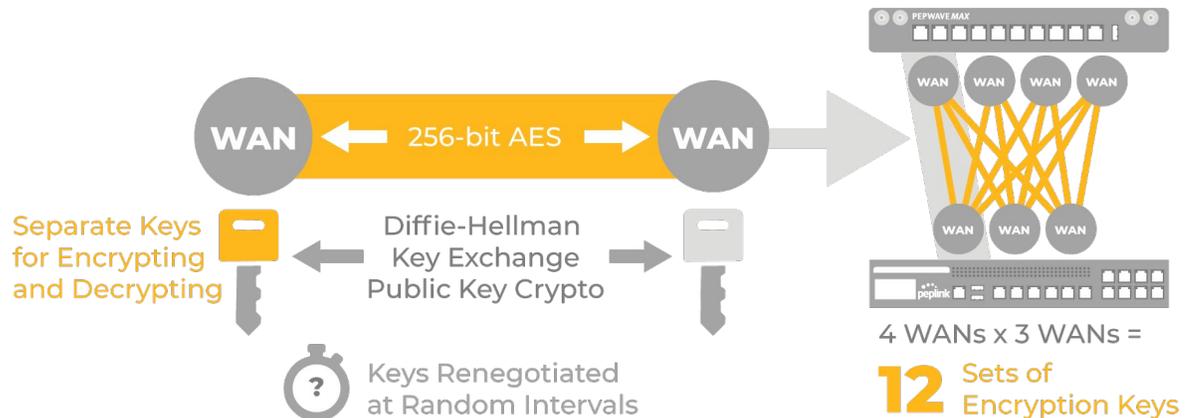
- 1) Navigate to PepVPN/SpeedFusion > Configuration
- 2) Select Network Topology
- 3) Select Participating Devices
- 4) Set Basic Profile Options



Introduction to SpeedFusion

SpeedFusion Security

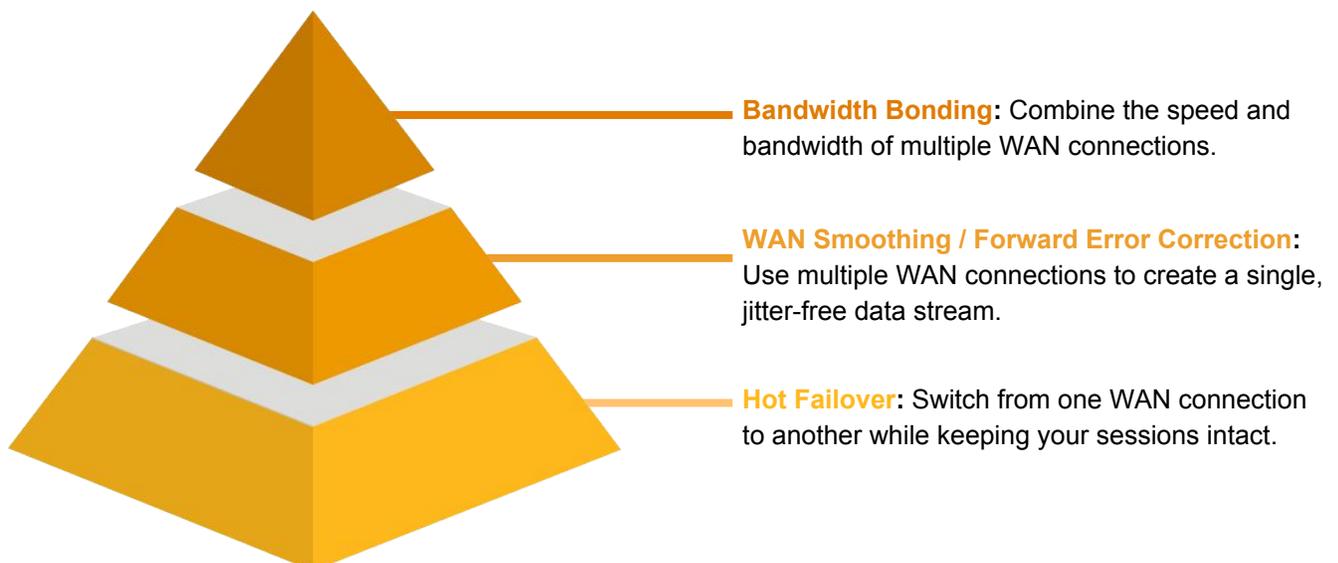
In SpeedFusion connection, a secure WAN-to-WAN link is established through Diffie-Hellman key exchange protocol that uses public key cryptography. Perfect forward secrecy is guaranteed by using DH exchanged master keys to derive keys for use in data encryption, and renegotiation of keys at random intervals. With the keys at hand, data passing through the WAN-to-WAN links are protected by AES cryptographic algorithm to provide 256-bit security level.



Once the bandwidth bonded SpeedFusion tunnel is formed, sessions are broken down and the resulting packets are sent separately across as available WAN-to-WAN connections. Because each WAN-to-WAN connection is encrypted separately, potential hackers would need to obtain the key to every WAN-to-WAN connection before they stand a chance of accessing the data contained within. This makes SpeedFusion SD-WAN technology virtually impervious to man-in-the-middle attacks.

Levels of SpeedFusion

From enterprise branch offices, to remote industrial sites, to vehicular and maritime deployments, SpeedFusion technology is adopted in multiple markets. To accommodate for different needs, we have developed different forms of SpeedFusion technology:

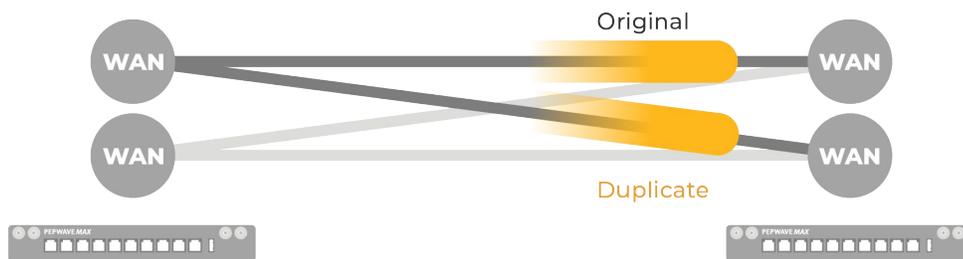


Hot Failover - Session Persistence

With this technology, two SpeedFusion peers can maintain session persistence even when switching between multiple WAN connections. The switching can occur between any WAN connection: fixed lines, cellular connections, or a combination of both. This technology enables authenticated and videoconferencing sessions to continue as connections change. Hot Failover serves as the core of all subsequent SpeedFusion technologies.

WAN Smoothing - Jitter-Free Streaming

For streaming technologies such as VoIP or videoconferencing, the connection quality strongly influences the user experience. High latency breaks the flow of conversations, while packet loss leads to jitter and missed sentences. WAN Smoothing reduces the impact of packet loss and improves latency in exchange for extra bandwidth consumption via packet duplication.



The amount of WAN Smoothing can be configured based on your deployment's bandwidth sensitivity:

- Normal - The total bandwidth consumption will be at most 2x of the original data traffic.
- Medium - The total bandwidth consumption will be at most 3x of the original data traffic.
- High - The total bandwidth consumption will be at most 4x of the original data traffic.
- Maximum - The total bandwidth consumption depends on the number of connected active WAN-to-WAN connection.

WAN Smoothing can also function with a single WAN-to-WAN link. In that case, the bandwidth consumption will be 2x the original data traffic.

Forward Error Correction - Lower Data Consumption

This technology has been developed to provide packet loss protection while minimizing bandwidth consumption. While WAN Smoothing duplicates actual packets, Forward Error Correction (FEC) sends additional reserve packets which could be used to mitigate the effects of packet loss via interpolation.

For each SpeedFusion profile, FEC can be configured as High or Low, with Low consuming 13.3% and High consuming 26.7% additional bandwidth.



Using SpeedFusion Behind a Firewall

By default, SpeedFusion uses **TCP port 32015** and **UDP port 4500** for establishing VPN connections and transmitting data. If IPsec or L2TP/IPsec services have been enabled, **UDP port 32015** is available as well. However, you can change the Data Port assignment in your SpeedFusion profile to another value.

Bandwidth Bonding

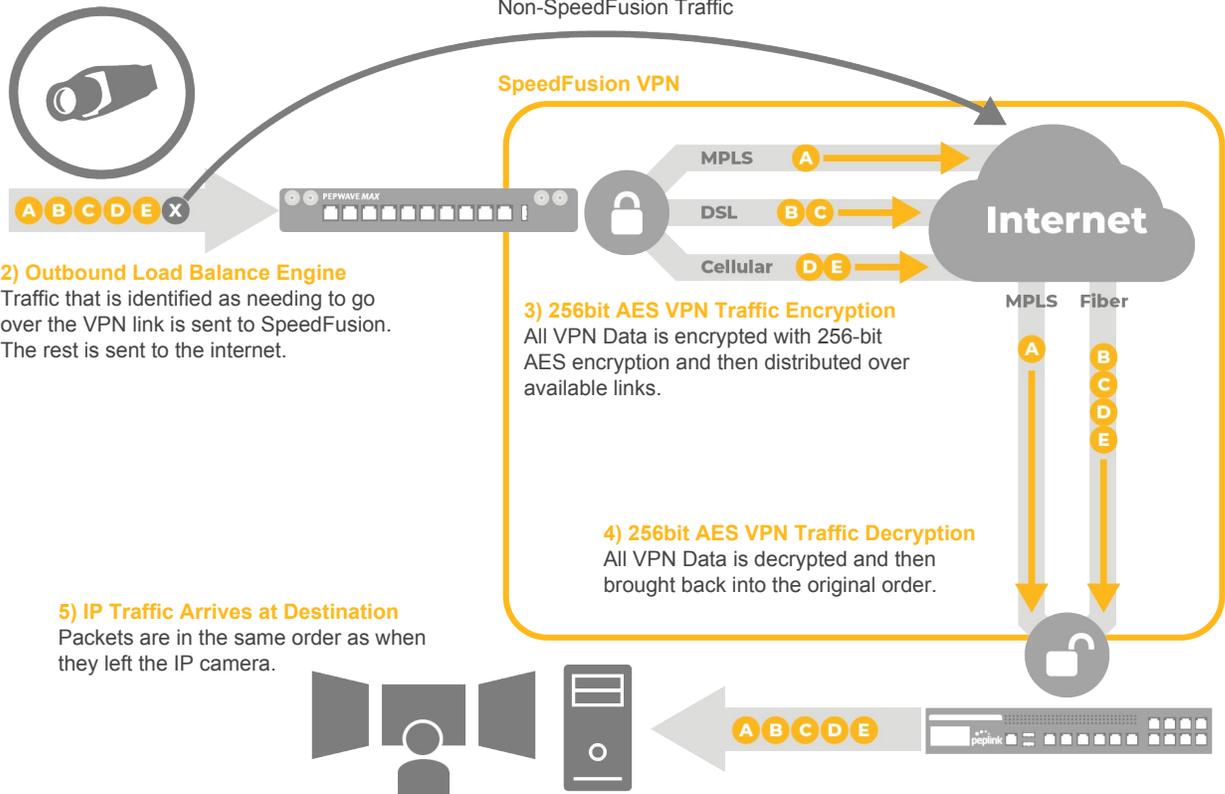
One of SpeedFusion's most powerful features is that it can use multiple WAN links to create a single logical VPN tunnel between devices. This allows SpeedFusion bonding to provide two key functions: bandwidth aggregation and VPN reliability.

All WAN links connected to a SpeedFusion-enabled device can simultaneously work together as a single logical VPN connection. If a link fails, SpeedFusion can detect this failure and seamlessly redirect traffic, at a packet level, across other available links. This WAN link failure detection at the packet level allows Bandwidth Bonding to provide VPNs with highly reliable and resilient site-to-site connectivity.

Below is an example of how SpeedFusion can be used to transmit IP video from a remote location to a central data center over multiple WAN Links .

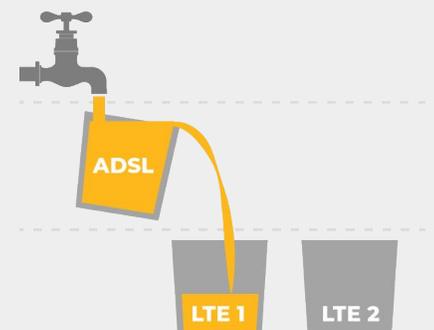
1) Data Packets

IP Data is sent from the IP CCTV camera over the LAN to the Remote router.



Traffic Overflow - Prioritize Affordable Links

Within a bandwidth-bonded SpeedFusion VPN tunnel, you could prioritize the order in which your links could be used. This is particularly useful for deployments where one link is more affordable than others. You can set the affordable WAN at first priority, and the more expensive links will be used only when the first is operating at capacity.



Bandwidth Bonding Overhead

SpeedFusion bonded VPN requires all transmitted data to be encapsulated in a special UDP stream. This stream contains additional packet headers with all the information needed to reconstruct the original data stream in the correct order at the remote location.

SpeedFusion adds an additional **80 bytes** of data to each packet sent over a SpeedFusion connection, no matter what size the original data packet is. This compares well to the 58 bytes of overhead required by IPsec, especially considering that SpeedFusion is not only providing advanced routing and load balancing but 256 bit AES encryption within the tunnel too.

SpeedFusion and The Internet Mix (IMIX)

Internet Mix (IMIX) is a measurement of typical internet traffic passing through network equipment, such as routers, switches, or firewalls. When measuring equipment performance using an IMIX of packets, performance is assumed to resemble what can be seen in real life.

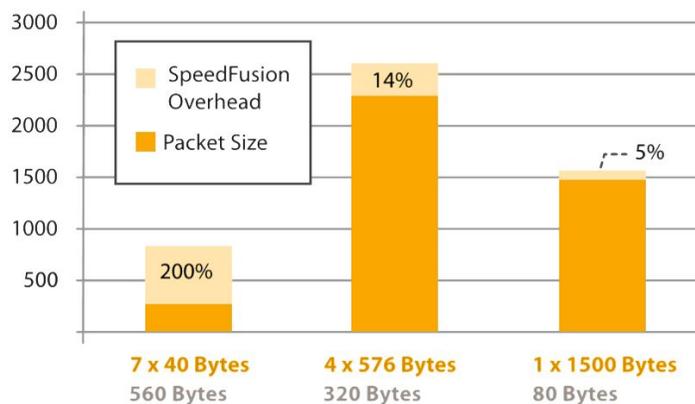
The IMIX traffic profile is used in the industry to simulate real-world traffic patterns and packet distributions. IMIX profiles are based on statistical sampling done on internet routers. More information about IMIX can be found here:

https://en.wikipedia.org/wiki/Internet_Mix

Packet Size	Packet #	Distribution in Packets	Bytes	Distribution
40	7	58%	280	7%
576	4	33%	2304	56%
1500	1	8%	1500	37%

A packet-level breakdown of the IMIX Standard (Total of 4084 Bytes)

The IMIX Standard and SpeedFusion Overhead



Effect of SpeedFusion Bandwidth Bonding overhead against the different packet sizes and quantities of the IMIX standard

As the chart on the left shows, when a SpeedFusion VPN tunnel is used to transmit IMIX data (4084 bytes), an additional 960 bytes of SpeedFusion overhead is required.

The SpeedFusion overhead is 19% of the total transmitted data (IMIX + overhead). Since it uses a fixed number of bytes per packet transmitted (an additional 80 bytes), SpeedFusion Bandwidth Bonding is much more efficient when transmitting larger packet sizes. At packet sizes of 1500 bytes, SpeedFusion adds just 5% bandwidth overhead, but at packet sizes of 40 bytes, SpeedFusion overhead rises to 200%.

SpeedFusion (80 Bytes)

$$\frac{7 \times 80 + 4 \times 80 + 1 \times 80}{(7 \times 80 + 4 \times 80 + 1 \times 80) + (280 + 2304 + 1500)} = 19\%$$

IPsec (58 bytes)

$$\frac{7 \times 58 + 4 \times 58 + 1 \times 58}{(7 \times 58 + 4 \times 58 + 1 \times 58) + (280 + 2304 + 1500)} = 14.6\%$$

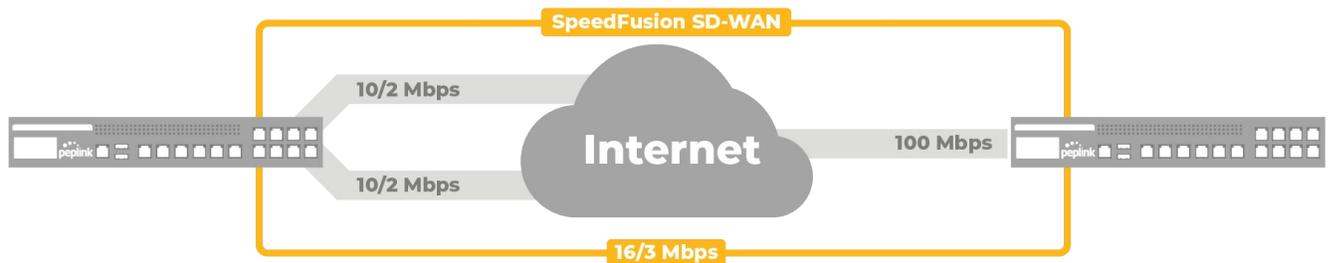
In Summary ...

For only 4.4% of additional overhead compared to IPsec, SpeedFusion includes bandwidth aggregation & WAN resilience.

Calculating Bonding Overhead

Using SpeedFusion for Bandwidth Aggregation Across Multiple Links

SpeedFusion is often deployed to new customers who have just a single internet connection but want more bandwidth. So, they buy an additional connection with the intention of using SpeedFusion to bond the two links. Without understanding SpeedFusion bandwidth overhead, a user might be confused by bandwidth availability using SpeedFusion bonding across links. Consider the following configuration:



Accounting for SpeedFusion bandwidth overhead and assuming the traffic passing across the links is similar to the IMIX standard mentioned previously, we can calculate available real-world bandwidth at the remote site:

$$\text{Download: } 10\text{Mb} + 10\text{Mb} = 20\text{Mbps} \times (1 - 19\%) = 16.2\text{Mbps}$$

$$\text{Upload : } 2\text{Mb} + 2\text{Mb} = 4\text{Mbps} \times (1 - 19\%) = 3.24\text{Mbps}$$

It is important to explain SpeedFusion bandwidth overhead to your end users so that they understand why they will not get full 20Mbps/4Mbps bandwidth when using VPN bonding. Remember, conventional VPN technology such as IPsec has an overhead of 14.6%. For only 4% of additional overhead, SpeedFusion provides **bandwidth aggregation & WAN resilience**.

Recommendation: Use SpeedFusion for Bandwidth Aggregation Across Links with Similar Capacities

For the highest performance, we recommend the use of WAN links with similar bandwidth profiles (within 50% of each other) and similar latency characteristics (we recommend within 150ms of each other). Use WAN links from different ISPs to allow for the best possible SpeedFusion throughput. Using at least two different ISPs offers the benefit of provider diversity, which means less chance of a technical (or even accounting/billing) error causing a network outage. Provider diversity also lessens the impact of bandwidth sharing, a common problem when using multiple circuits from a single provider.

MPLS Complement / Alternative

Customers with existing expensive MPLS & Leased Line circuits can take advantage of SpeedFusion's ability to build Hybrid WAN, complementing or completely replacing these circuits. A SpeedFusion built Hybrid WAN:

- 1) Provides SP/Carrier Diversity & Redundancy
- 2) Lower cost per Mbps

More Than Bandwidth Aggregation

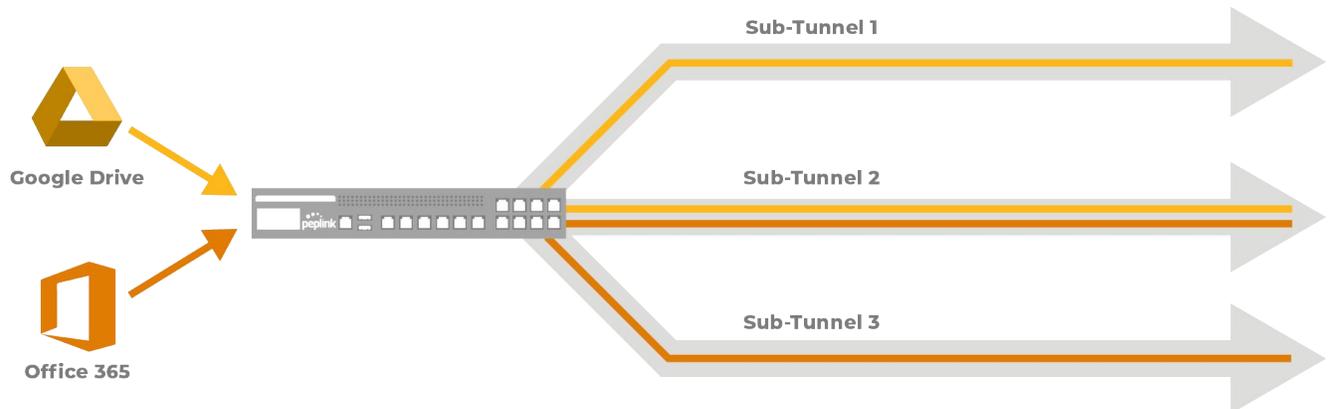
Customers can take advantage of SpeedFusion's reliability and use low-cost DSL circuits to achieve higher reliability and throughput than comparable private circuits.

Up to 80% Less cost.

SpeedFusion Outbound Policy

SpeedFusion Traffic Steering

Modern networks carry a variety of traffic such as video streams, ERP sessions, HTTP sessions, and more. Depending on your deployment, you may need to give priority to certain kinds of traffic while throttling or even blocking others. SpeedFusion gives you granular control of how different kinds of traffic travels within your network.

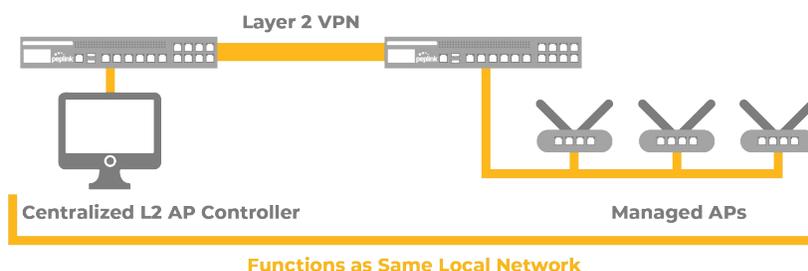


This is achieved through user defined sub-tunnels. Within a SpeedFusion connection, you can define up to five sub-tunnels. Within each sub-tunnel, you can specify which WAN connections to use, set bandwidth limits, as well as configure WAN smoothing and FEC. Once you have defined these tunnels, you can then use outbound policy to steer traffic based on protocol (TCP vs UDP), Port, and Application.

Deep Packet Inspection: Real-Time Application-Level Packet Recognition

Peplink has support for Deep Packet Inspection (DPI), which enables routers to scan passing packets and recognize traffic from various applications (such as Skype, Google Hangouts, Spotify, Dropbox, and BitTorrent). This information can then be used to define firewall rules, define outbound policy, and assign SpeedFusion Sub-Tunnels.

Multi-Site Layer 2 VPN

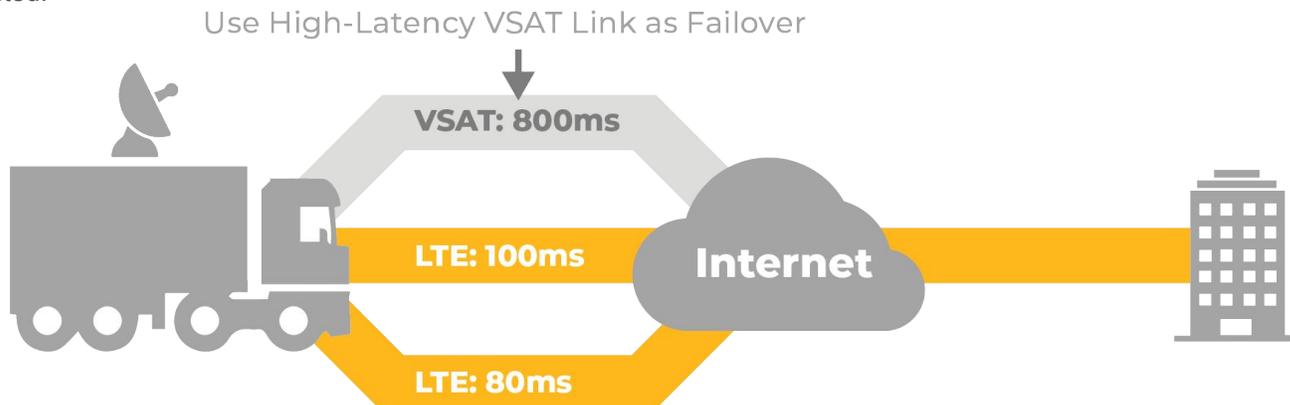


By forming a SpeedFusion Layer 2 VPN, you can get devices from different sites to operate as if they were on the same local network. This technology is useful for remotely operating devices such as printers and access points, sharing network resources such as NAS drives, as well as facilitating server-to-server communications.

A Layer 2 VPN can be applied to the VLAN trunk or a specific VLAN. If it is applied to the trunk, then all traffic from the router will go through the VPN and will be considered as LAN traffic by the remote site. Alternatively, if you define the Layer 2 VPN on a specific VLAN, then only traffic from that VLAN will be considered local traffic by the remote site.

Addressing Latency and Packet Loss

When latency characteristics are the same across connected WAN links, it has very little effect on SpeedFusion bandwidth throughput. However, when the latency of WAN links varies considerably, bandwidth throughput is affected.



Packet Loss in High Latency Environments

In the example above, there is an LTE connection which is highly susceptible to packet loss. Because the latency across the SpeedFusion link is equalized to the highest latency link (800ms), if a packet is lost on the LTE link, SpeedFusion will take longer to spot the packet loss (800ms+).

In certain conditions, such as regular timed packet loss on the above LTE link, combined with typically high latency found on these types of links, the TCP protocol method of retransmitting lost packets can have a drastic effect on available bandwidth over the VPN. This is another reason why we recommend that, whenever possible, high latency links be used for failover and not as an active SpeedFusion WAN link.

Managing High Latency Links

In cases of high variation in WAN link latency, the best approach (assuming there is enough bandwidth on low latency links) is to allocate lower latency links for SpeedFusion while setting higher latency links as failover connections.

Another approach is to use higher latency links for specific direct traffic types that are not as latency sensitive (like direct internet access) while reserving lower latency links for other important corporate traffic that needs to transit the VPN, such as VoIP and ERP.

Note: *Using UDP traffic over SpeedFusion can provide higher throughput than TCP which has restrictive flow control.*

Mitigate with SpeedFusion WAN Smoothing and FEC

Peplink has two purpose-built technologies to mitigate the effects of packet loss and minimize latency:

WAN Smoothing duplicates packets across multiple WAN connections. This technology produces the lowest latency and the minimum packet loss in exchange for bandwidth consumption. This technology is best for two-way communications such as VoIP and videoconferencing.

Forward Error Correction (FEC) uses algorithms to interpolate lost packets. FEC consumes less overheads than WAN Smoothing in exchange for latency. This technology is best used for one-way data traffic such as video streaming.

WAN Link Testing - Use speedtest.net or nperf.com

Whatever WAN connections you are using, it is always a good idea to test each individually and repeatedly to discover its maximum throughput in both directions. Remember, bandwidth availability can vary throughout the day, especially if using cellular or fixed lines with variable contention.

Customers are often surprised by how bandwidth availability can vary between links from the same ISP and how different actual bandwidth availability can be from that advertised by the ISP.

Cellular and Satellite Bandwidth Availability



The amount of bandwidth available on a LTE or satellite data connection is dependent on a number of factors:

- **Signal Strength** -Determined by the distance to the nearest cellular tower (or visibility of the satellite) and the subsequent signal quality received.
- **Backhaul Bandwidth Availability**- From the cellular tower to the ISP's core network or from the satellite ground station to the ISP's core network.
- **Device Contention** -At the tower or satellite you are connected to (determined by the number of active subscribers on a tower or satellite at any given moment).

We frequently see users who have 100% signal strength yet get only a small percentage of the bandwidth. This can be because there simply isn't enough backhaul bandwidth at a cell tower or ground station to the ISP's backbone, and what is available is saturated. When it is a matter of contention of the service, the easiest way to improve your chances of getting more bandwidth is to create another connection using an additional vSAT or cellular modem connection. This allows you to get twice as much of the time allocated to your devices, and the amount of bandwidth shared with you for your SpeedFusion connection is increased

Fixed Line Contention

Most internet connections are provided as a contended service. This means that although your provider has advised you will get up to 24Mbps broadband over DSL for example, depending on how oversubscribed your DSL service is (literally how many people in your area are connected to the ISP's service), the bandwidth that's actually available at any given moment could be considerably less.

The amount of bandwidth available on a contended service can vary considerably over the period of a day (and even minute to minute), with higher speeds possible during working hours compared to the evenings when your neighbours are home and using the same internet service heavily.

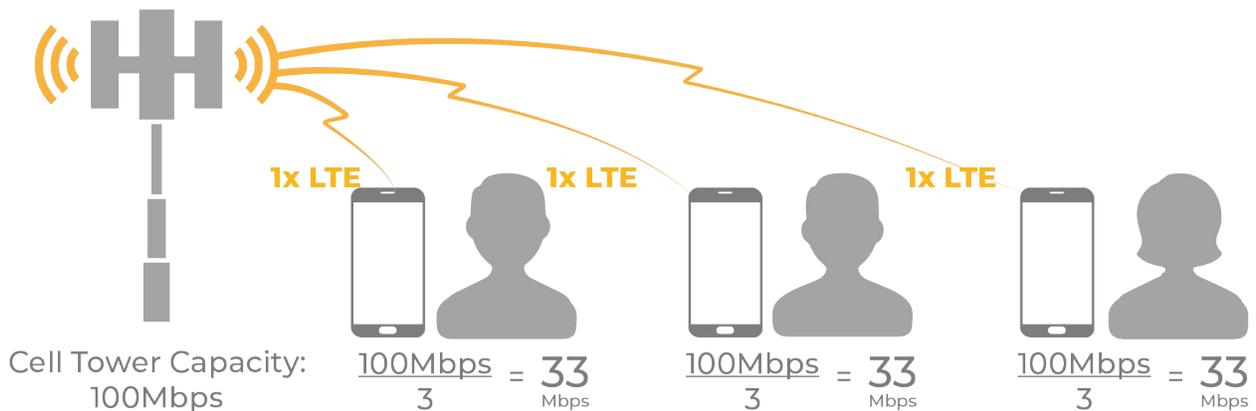
Adding an additional fixed line service from the same ISP can give you a 'bigger share' of the bandwidth that's available. 1 :1 contended services are available from ISPs to counter this issue but are naturally more expensive than 20:1 or 50:1 services.

Application Considerations

The Benefits of Using Multiple LTE Connections on Contended Cell Towers From the Same Mobile Operator

Verizon and other LTE providers use a process called windowing/time-slicing when multiple subscribers connect to their LTE services.

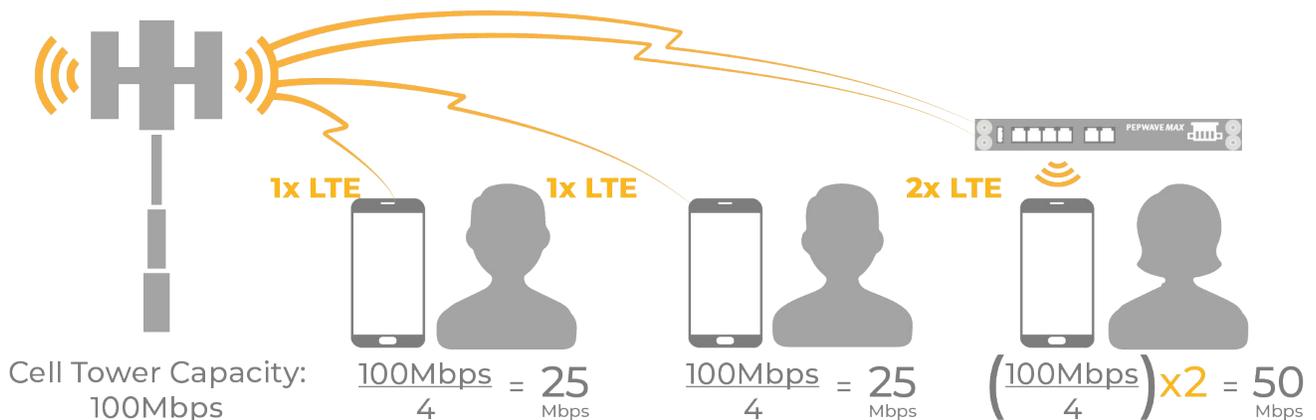
Single Cellular Connection



In this example, three users with three LTE-enabled devices connect to a cell tower. Each gets 33Mbps, which is a third of the available bandwidth at the tower.

$$100\text{Mbps} / 3 = 33\text{Mbps}$$

Dual Cellular Connection



In this example, the third user is using a Pepwave MAX device and has installed an additional LTE data SIM. There are now three devices with four LTE connections connected to the cell tower.

$$(100\text{Mbps} / 4) \times 2 = 50\text{Mbps}$$

The first two users get one quarter of the available bandwidth (25Mbps), but the third user gets two quarters (or half) of the available bandwidth using his two LTE connections.



<https://forum.peplink.com>



<https://www.youtube.com/PEPLINK>



<https://www.facebook.com/peplink/>



<https://twitter.com/peplink>



<https://www.linkedin.com/company/peplink>