HP Velocity

Fifth Edition: August 2016

First Edition: June 2012

Document Part Number: 689165-005

Quality of Experience with HP Velocity

HP Velocity is a Quality of Service (QoS) software engine that improves the overall Quality of Experience (QoE) for real-time network applications.

Real-time streaming applications delivered over data networks can be compromised by packet loss, transmission latency, and Wi-Fi congestion. This results in stop-and-go behavior, loss of interactivity, and an overall reduction in the throughput of an application.

HP Velocity easily integrates with existing systems to protect application flows between HP thin clients and HP Velocity-enabled servers, including terminal services and virtual desktops hosted by VMware®, Citrix®, or Microsoft® hypervisors. It improves the QoE of a streaming application by addressing the underlying problems found in today's networks: packet loss, latency, and congestion.

HP Velocity automatically ensures the best end-user QoE by optimizing the use of the underlying network to meet the requirements of bandwidth-intensive applications over wired and wireless, managed and unmanaged networks. By managing packet loss, Wi-Fi congestion, and the impact of network latency, HP Velocity improves QoS, providing a better experience for the end user.

HP Velocity continuously monitors end-to-end network conditions to select the most appropriate data-delivery mechanism. Rich reporting capabilities enable HP Velocity to provide end-to-end visibility to key flow information. The system's network monitoring and reporting tools capture in-depth statistics to log files, making them available for analysis and troubleshooting.

Benefits of HP Velocity

Adaptive network analysis	HP Velocity continuously monitors end-to-end network conditions for individual data flows, providing adaptive optimizations and application-flow protection.
Monitoring	HP Velocity collects and reports an extensive set of statistics: • System information • Endpoint and per-flow network statistics: network-loss rates, corrected-loss rates, throughput, and latency
Packet loss protection	HP Velocity provides packet-loss protection, which accelerates UDP application flows by reducing packet loss and the number of retransmissions.
Latency mitigation	HP Velocity boosts application QoE in a high-latency environment. By actively adapting TCP, HP Velocity automatically calibrates congestion control parameters for each TCP flow, based on the conditions present in the network.
Congestion detection	HP Velocity automatically detects network congestion and adapts accordingly to maximize QoE.
Wi-Fi optimization	HP Velocity automatically reduces latency and transmission times for wireless networks and minimizes protocol overhead, resulting in improved QoE for Wi-Fi applications.
Seamless integration	HP Velocity is a transparent, plug-and-play solution that provides QoE benefits to applications and end users.
Lightweight	HP Velocity is delivered as a lightweight implementation, which achieves QoE benefits while keeping system resource use to a minimum.

HP Velocity components

HP Velocity consists of the following components, which form an integrated system designed to improve overall QoE.

|	
Policy Engine	• Matches application requirements • Tunable through API or the Group Policy feature in Windows®
Network Profiler	• Profiles the network path between HP Velocity-enabled endpoints • Identifies the type of network connectivity (wired or wireless) and measures key network metrics (packet loss, latency, and bandwidth constraints) independently for each flow • Continually updates the QoS Controller with current end-to-end network conditions
QoS Controller	• Uses the current and trending network conditions provided by the Network Profiler to intelligently activate and adjust the optimizers
Optimizers	• LiveQ: Provides end-to-end packet-loss protection to UDP application flows by automatically adapting the amount of added redundancy. For more information, see LiveQ on page 3. • LiveWiFi: Improves performance for congested Wi-Fi networks and increases the efficiency of TCP over Wi-Fi. LiveWiFi protects application flows by leveraging Wi-Fi standards to minimize latency, resulting in the prioritization of HP Velocity-protected flows. For more information, see LiveWiFi on page 4. • LiveTCP: Optimizes TCP throughput over all networks and provides latency mitigation for the RDP, HP Remote Graphics Software (RGS), PCoIP, VMware Blast, and Citrix ICA protocols. LiveTCP improves the throughput of streaming and remote desktop applications by modifying TCP flow control mechanisms to perform better in high-latency environments. For more information, see LiveTCP latency mitigation on page 6. **Note:** HP Remote Graphics Software (RGS) versions 6.0 and newer include HP Velocity as a built-in feature that is integrated into the HP RGS user interface on both the server side (RGS Sender) and client side (RGS Receiver). If you want to utilize HP Velocity solely for HP RGS, you do not need to install any standalone HP Velocity software components. For more information about using HP Velocity in HP RGS, see the User Guide for your version of HP RGS. HP RGS versions prior to 6.0 are not compatible with and do not include HP Velocity.

Figure 1: HP Velocity components
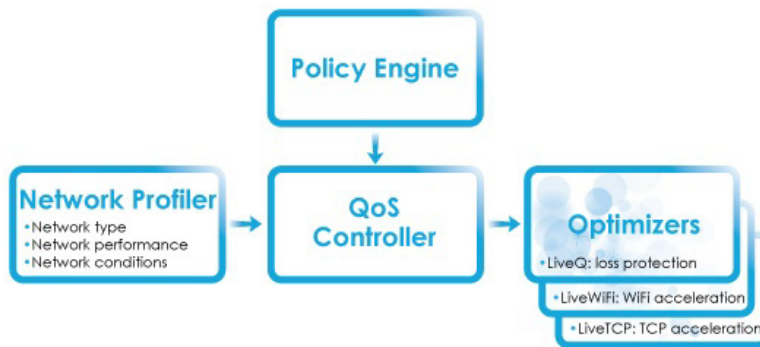
LiveQ

Packet loss occurs when packets traversing a network do not reach their destination. A number of factors contribute to packet loss, including signal degradation over the network, network congestion, corrupted packets, and hardware issues.

In UDP-based transport, which does not guarantee packet delivery, even a small amount of packet loss can cause applications to slow down or reduce the bit rate. This results in digital artifacts, smeared video, broken and choppy audio, or low frame rate in applications such as PCoIP, RDP, VP8, SIP, VoIP, and video.

LiveQ outputs the original number of input packets as well as extra packets containing redundancy information. The number of packets that are used to encode the redundant packets and the number of redundant packets generated varies based on current network-loss conditions and the configured Target Loss Rate.

Figure 2: LiveQ transformation of packets
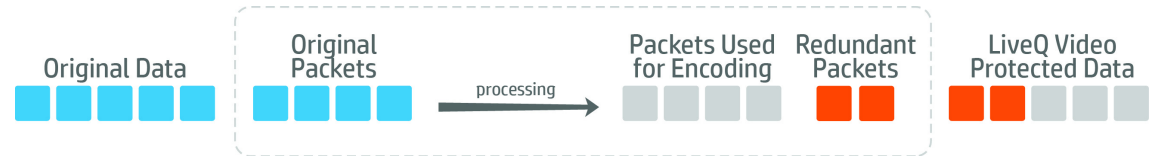
Target loss rate

Different applications tolerate different levels of packet loss. Some applications perform poorly with a small amount of packet loss while other applications perform satisfactorily even with significant amounts of packet loss. HP Velocity adjusts its operation to ensure that each application is protected from experiencing too much packet loss.

The target loss rate (TLR) is the amount of loss that an application can tolerate while still delivering an acceptable QoE. The recommended TLR for thin-client applications is 0.04%.

Encoding modes

The encoding modes determine the amount of redundancy added to an application flow. HP Velocity automatically adapts the encoding mode for a flow according to the configured TLR and current network conditions.

The values in Table 1 show how HP Velocity adapts the LiveQ encoding mode to keep the corrected loss level below a TLR of 0.04% for an application flow with a bit rate higher than 1 Mb/s.

**Table 1: LiveQ encoding**

Network Loss (%)	Redundancy (%)
0.04 - 0.8	10
0.8 - 1	20
1 - 1.8	30
1.8 - 2.8	40
2.8 - 3.9	50
3.9 - 5.9	60
5.9 - 6.2	70
6.2 - 7.9	80
7.9 - 8.8	90
8.8+	100

LiveWiFi

Wi-Fi is constrained for a number of technical reasons and, as a result, delivers a much lower network performance experience than wired networks. High-end, consumer-grade Wi-Fi access points (APs) advertise that they are capable of speeds of 300 to 450 Mb/s. In reality, Wi-Fi capacity is much lower than advertised.

- Distance (signal strength): The further away a Wi-Fi-enabled device is from the AP, the lower the signal strength, which in turn lowers the effective available throughput.

- Noise from other devices: The most common frequency used for Wi-Fi is the 2.4 GHz band, which is also used for household devices, such as cordless phones, baby monitors, and microwave ovens. Throughput is degraded if any of these devices are used within range of a Wi-Fi network.

- Noise from other APs: There are only three non-overlapping channels that can be used by 2.4 GHz Wi-Fi. In areas with a high density of APs, such as a city neighborhood or an office tower, it is likely that there are several other networks using the same channel, leading to speed degradation.

- Older devices: If an 802.11n-enabled AP tries to connect to an older device that supports only 802.11b, the AP must drop the supported speeds for everyone while the 802.11b device is operating. This significantly reduces the throughput that 802.11n devices can achieve.

- Management overhead: Wi-Fi networks carry a significant amount of management overhead. As APs advertise their presence, wireless clients must regularly probe the AP to notify it that they are still there. When there are a large number of Wi-Fi devices present in a single location, the management overhead becomes excessive and the capacity of the Wi-Fi network is reduced.

- Acknowledgments: All Wi-Fi packets must be acknowledged by the receiver to ensure successful delivery. These acknowledgment packets are sent by a Wi-Fi device every time it receives a packet. When combined with the overhead of protocols like TCP, this can result in three out of every four Wi-Fi packets becoming overhead packets, with only one out of four packets containing useful data.

Figure 3: Wi-Fi overhead in a non-optimized flow
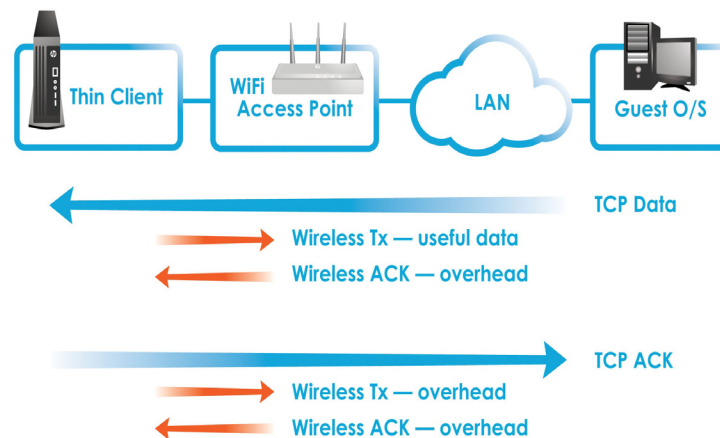
- Half duplex technology: A wired Ethernet network is full duplex, meaning that a device can send and receive, or upload and download, simultaneously. Wi-Fi is half duplex; thus, if a client sends data to the AP, the AP cannot send data to the same or any other client at the same time. For two-way communication, which includes most applications used over the Internet, such as video or voice chat, this essentially halves throughput when compared to a full-duplex technology.

- Retransmissions: Not all Wi-Fi packets are successfully received the first time they are sent. A client might not receive a packet because of collisions or insufficient signal strength. A single bit error in a packet results in the retransmission of the entire Wi-Fi packet, which might happen repeatedly for the same Wi-Fi packet.

- Wireless contention: When a client wants to send data and the channel is occupied, the client must wait; otherwise, collisions will occur and the data will be corrupted. Once the channel clears, the client must wait even longer before it can attempt to send. Similar to trying to cross a single-lane bridge, if there are cars wanting to cross from both directions, everyone must wait their turn and there is uncertainty about whose turn is next.

  If a user on a wireless network is using large amounts of bandwidth through activities such streaming a video or downloading a file, it is more difficult and takes longer for another device to send, even when performing a task that requires very little data, such as checking email.

LiveWiFi addresses these issues using Wi-Fi prioritization and TCP over Wi-Fi.

Wi-Fi prioritization

LiveWiFi gives HP Velocity traffic a higher priority than other traffic on the network. Both endpoints mark HP Velocity traffic to prioritize packets moving in either direction. LiveWiFi leverages the Wi-Fi Multimedia (WMM) Standard to reduce packet wait times compared to other traffic in the network. This results in lower latency, less jitter, and higher throughput.

TCP over Wi-Fi

Standard TCP traffic acknowledgments interfere with the transmission of useful data on a Wi-Fi network. With LiveWiFi enabled, HP Velocity modifies the rate of TCP acknowledgments to reduce protocol overhead and improve network performance. This can liberate up to 15% of additional usable Wi-Fi bandwidth.

LiveTCP latency mitigation

Real-world IP networks introduce latency and packet loss to application flows. A primary cause of these issues is network congestion.

Latency and packet loss degrade the end-user QoE, especially for applications that use TCP for reliable data transmission, such as streaming video, chat, remote desktop, and file transfers. This results in unresponsive user interfaces for remote desktop users, extended wait times for file transfers, and choppy video with reduced frame rates.

TCP congestion avoidance mechanisms

TCP is designed to send data as rapidly as possible until it detects packet loss. Once packet loss is detected, TCP halts data delivery and reduces transmission speed while lost packets are retransmitted. Built-in algorithms enable TCP to avoid excessive congestion of the network. As packet loss dissipates, TCP increases the speed of data delivery.

TCP algorithms

TCP can be controlled by different algorithms to accommodate different types of networks. However, only one network type at a time can be accommodated under native TCP. For example, LANs are characterized by low packet loss and high bandwidth. Remote office networks typically have high packet loss and high latency. Wi-Fi networks have high jitter (that is, the latency is highly variable) and variable throughput capacity.

TCP tuning

Most operating systems allow only a single TCP algorithm to be active with a single set of settings. If a network is optimized for LAN traffic, WAN connections are less efficient. While the TCP stack can be tuned, the effort required by network administrators is a deterrent. LiveTCP automatically provides tuning, which boosts TCP throughput over all networks and provides latency mitigation for streaming applications.

LiveTCP congestion control

LiveTCP adapts its algorithm to automatically provide congestion control that is optimized independently for each TCP flow (LAN, remote network, Wi-Fi) and the conditions in the network. LiveTCP achieves this without requiring intervention from IT.

LiveTCP improves on native TCP by accelerating the speed at which thin-client protocols transmit data. For example, for RDP, which is highly sensitive to latency, LiveTCP can improve transmission speeds up to 10 times; for HP RGS, two to three times; for VMware Blast, up to two times; and for ICA, two times.

In public networks where congestion is often a problem, such as at airports, hotels, and coffee shops, LiveTCP improves overall throughput by more effectively managing shared bandwidth.

LiveTCP is configured to provide latency mitigation for application flows between two HP Velocity-enabled endpoints (that is, in a peering configuration) or from an HP Velocity-enabled endpoint to a non HP Velocity-enabled endpoint (single-sided). In both configurations (Figure 4 and Figure 5), the optimized LiveTCP algorithm takes precedence over the native TCP algorithm.

Figure 4: LiveTCP algorithm in a peering configuration
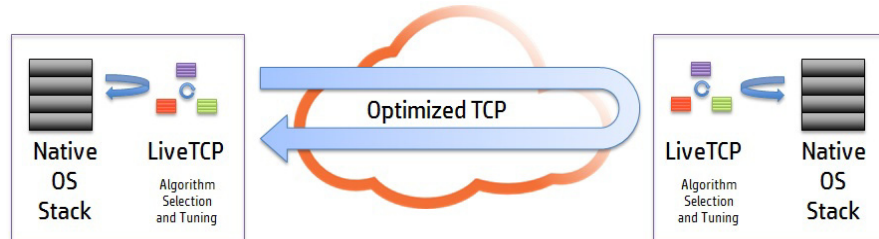
Figure 5: LiveTCP algorithm in a single-sided configuration
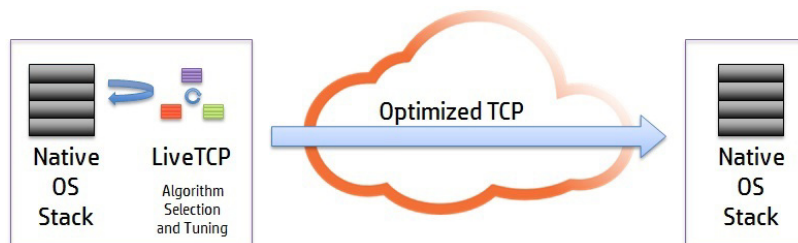
Summary

HP Velocity patented technology is available exclusively on HP thin clients. It provides a seamless, integrated QoE solution for both managed and unmanaged networks. This allows organizations to gain valuable insight into their thin client application flows, offload more expensive infrastructures (such as MPLS networks) to lower-cost networks, and achieve maximum performance regardless of network conditions.

The major benefits of HP Velocity include the following:

• End-to-end visibility of thin client application flows

• In-depth monitoring and reporting

• Packet-loss protection

• Wi-Fi acceleration

• Latency mitigation

• Support for all VDI protocols, including RDP, HP RGS (version 6.0 and later), PCoIP, VMware Blast, and ICA

HP Velocity protects application flows between HP thin clients and HP Velocity-enabled servers, including terminal services and virtual desktops hosted by VMware, Citrix, or Microsoft hypervisors.

To obtain HP Velocity and associated documentation, go to http://www.hp.com/go/velocity.