

# The Different Layers of WAN Deduplication

Understanding the Differences Between TCP and IP-based  
WAN Deduplication

# The Different Layers of WAN Deduplication

## Understanding the Differences Between TCP and IP-based WAN Deduplication

On the surface, different vendors' Wide Area Network (WAN) deduplication solutions sound a lot alike. When one digs deeper, however, it becomes apparent that there are fundamental architectural differences that have a tangible impact on the performance of different applications across the WAN.

In a previous paper, entitled, *The Benefits of Byte-Level WAN Deduplication*, we addressed one major difference—the usage of tokens vs instructions when referencing repetitive patterns within a common data set. This affects the level of granularity provided by the deduplication algorithms and thus the level of performance offered in specific application environments, especially those with dynamic data.

In this paper, we will address another major difference—the layer at which deduplication takes place. More specifically, we will explore the differences between TCP and IP layer deduplication solutions, and how they translate into real-world advantages and disadvantages. As one will see, choosing the right architecture can have an enormous impact on the scalability of a WAN acceleration solution in terms of the number of sessions supported and the types of applications that can be optimized.

### TCP vs. IP Layer Deduplication

By their very nature, WAN deduplication solutions must be deployed in networks such that they are directly in the path of all data transmitted to and received from the WAN. WAN deduplication solutions are typically implemented as network appliances and deployed either physically or virtually in the data path. This enables these solutions to intercept data, remove as much repetitive information as possible, and then transmit the reduced data onto the WAN. Removal of repetitive data is accomplished whenever the appliances are able to reference previously transmitted data.

Theoretically, interception of network data can be performed at any layer of the protocol stack. Practically, though, two broad classes of implementations have evolved amongst WAN deduplication solutions—those which operate at the TCP layer (Layer 4 of the OSI model—also known as “transport layer”) and those which operate at the IP layer (Layer 3, or “network layer”).

TCP layer deduplication solutions operate by terminating (or “proxying”) TCP sessions that flow through the WAN optimization appliance. TCP packets received on the LAN side of the appliance are queued up for reassembly into segments by the appliance's TCP stack. After termination and reassembly, segments from each TCP flow are analyzed by the deduplication algorithm.

In contrast, IP layer deduplication solutions operate at the network (or IP) layer of the protocol stack. As soon as a packet is received on the LAN side of an appliance, the headers are removed and the packet is delivered to the deduplication algorithm for analysis.

While the differences between TCP and IP based deduplication schemes seem subtle, the impact of using one approach versus the other can be quite significant. This is particularly true in the following three areas:

- **Applications Supported.** TCP-based schemes only work on applications that use the TCP protocol for transport. This limitation renders such implementations ineffective for UDP applications, such as voice, video, and certain replication applications (for example, Veritas Volume Replicator, Aspera, Isilon, and CLARiiON Disk Library).

---

*While the differences between TCP and IP-based deduplication schemes seem subtle, the impact of using one approach versus the other can be quite significant.*

In addition, TCP-based schemes do not work with applications that run natively over IP like certain Fibre Channel over IP (FCIP) implementations, anything encapsulated in IP (GRE, IP-IP, IPv6 tunneled in IPv4, etc.) and proprietary IP-based applications.

In contrast, IP-based schemes will work on all IP traffic, regardless of transport protocol.

- **Flow Capacity.** TCP establishes a flow or connection whenever information is transferred between two end points. In the typical enterprise network, there are approximately 10 to 50 flows open at any given time per employee, resulting in the potential for hundreds of thousands of connections in large networks.

TCP-layer implementations often have difficulties handling large volumes of flows because they typically rely on the underlying operating system to perform TCP reassembly and resegmentation operations. In this type of implementation each TCP-proxied flow requires two operating system sockets — one for the LAN side and one for the WAN side connection. While operating systems like Linux are designed to efficiently terminate hundreds of connections/sockets, they are not optimized to simultaneously support tens or hundreds of thousands of connections/sockets.

To address this issue, many TCP-based solutions will limit the number of TCP flows supported per appliance, limiting the scalability of these solutions. Even with these limits, it is very difficult for these socket-based systems to provide fair service to every single flow.

In contrast, IP-based solutions decouple TCP proxy functions from the underlying operating system, enabling them to support hundreds of thousands of simultaneous flows per appliance. This eliminates any flow capacity bottlenecks that otherwise might be seen in WAN optimization devices.

- **Latency.** TCP-layer implementations often add significant latency when doing deduplication. There are two primary reasons for this. The first is related to the way the appliance’s operating system schedules service when there are hundreds or thousands of active connections. More specifically, the operating system usually has no way to prioritize per-socket service in a predictable way, preventing it from placing latency-sensitive traffic ahead of other traffic in the deduplication queue.

Secondly, many TCP layer schemes use a token-based approach to deduplication, whereby data is buffered into large groups before pattern matching takes place. This allows for great data reduction when dealing with bulk data transfers, but it adds significant latency, which can be problematic when dealing with real-time and interactive traffic like Citrix and Remote Desktop Protocol (RDP). With the proliferation of multimedia (VoIP, video, and real-time data) and the widespread usage of virtualized desktop applications (that is, thin clients), this is becoming an increasing concern amongst enterprises.

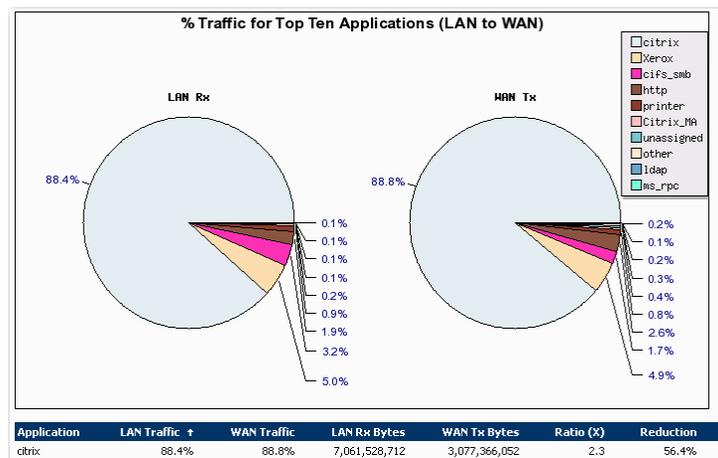


Figure 1: Silver Peak adds very little latency when performing deduplication, enabling data reduction on real-time traffic like Citrix.

- IP layer deduplication solutions have better prioritization capabilities because TCP services are decoupled from the operating system. In addition, they often use an instruction-based approach to deduplication instead of tokens, which perform pattern matches in real-time as every byte of data enters the appliance. This avoids latency that comes from buffering, making it ideal for time-sensitive applications.

## Supporting Latency-Sensitive Applications Across the WAN

Thin client solutions (such as Citrix, Microsoft's RDP and VNC) allow the centralization of enterprise applications, which enables lower capital and operational expenditures. Unfortunately, though, thin clients are highly sensitive to latency. Any increase in latency results in delayed screen refreshes and delayed keystroke "echoes". This leads to frustrated users who are far less productive than they should be or who refuse to utilize the thin-client applications altogether, resulting in increased costs for the enterprise.

A similar challenge exists when trying to support multimedia solutions across the WAN, such as video streaming. When transmitting this type of traffic, any increases in latency will cause frames to be dropped or delivered out of order. The result is "jerky" performance.

Thin client applications, multimedia solutions, and latency-sensitive applications can all be classified as real-time in nature, meaning that any increase in latency significantly degrades the applications, often to the point where they are unusable.

TCP buffer techniques severely degrade the performance of real-time applications, often to the point that they perform worse than without any WAN optimization. Consequently, TCP layer WAN acceleration solutions typically "bypass" (that is, don't try to optimize) latency-sensitive applications. (As a good rule of thumb, check the vendors' user manuals to see what applications are bypassed by default....)

## Using IP Layer Deduplication to Optimize Real-Time Traffic

The question often arises — can you actually deduplicate real-time traffic? Some applications, like VoIP, contain many packets that are fairly unique in nature, which makes deduplication difficult unless it is voicemail or a recorded message. In these instances, the traffic is best optimized using other WAN optimization techniques, like loss mitigation, QoS, and header compression. Other real-time applications, like Citrix and video streaming, can benefit enormously from WAN deduplication under the right circumstances — that is, when content is repetitive over time. Figure 1 is a screenshot showing 56% data reduction at a global law firm using Citrix ICA presentation server.

Because Silver Peak does not bypass any traffic, a full array of WAN optimization capabilities can be applied to all real-time traffic. In some instances, deduplication will play a major role. In other instances, Forward Error Correction (FEC), Packet Order Correction (POC), Quality of Service (QoS), and other techniques will be the primary drivers for improved application performance. Since it all depends on the type of traffic being sent, enterprises require an IP-based solution to give them the flexibility they need to adapt to different application requirements.

## Conclusion

Silver Peak's Network Memory works at Layer 3 of the OSI model, avoiding the scalability and performance issues associated with TCP-based deduplication schemes, as described above. The result is real-time optimization for all IP traffic, including Citrix, video streaming, data replication, and other latency-sensitive applications. By working at the network layer, Silver Peak also avoids TCP flow limitations that can plague some WAN acceleration appliances, while ensuring that a full range of optimization techniques can be applied to everything that goes across the WAN.

By providing an IP-based solution for WAN optimization, Silver Peak delivers the best performance, across the most applications. As a result, enterprises get the best return on their WAN acceleration investment.