

How Riverbed's Data Replication and Disaster Recovery can Assist with Compliance Goals



Companies are experiencing huge upticks in distributed data creation with rapidly increasing interconnected devices and nodes from large numbers of applications and locations. The expansion of computing and analytics capabilities drives the need to obtain greater value from all data sources. Getting that data to the correct location and into the hands of people who need it quickly is a considerable challenge.

This requires moving data from one location to multiple sites in a timely manner without impacting normal workflow and business operations. These issues become more troublesome as data, distance, location, latency, and bandwidth increase. Companies are faced with unpredictable, slow, and costly solutions, with shared data becoming stale in transit and arriving later than needed. Network limitations impact this behavior tremendously – ultimately affecting productivity and profits.

Enter Riverbed SteelHead

One of our Acceleration solutions, Riverbed SteelHead, enables companies to efficiently distribute data – exceptionally fast and to scale – across multiple clouds and distributed and hybrid networks. This allows organizations to achieve efficient Recovery Time Objective (RTO) and Recovery Point Objective (RPO), whether you're mirroring databases, backing up desktops and servers or backing up hot standby sites and repositories.

SteelHead rapidly moves data from edge locations, data centers, and cloud providers – all securely, with market-leading encryption for your complete peace of mind. Our solution is agile and can be used as a hardware or virtual form factor. What's more, it comes in multiple sizes to fit your enterprise. SteelHead also complements other vendor solutions like SD-WAN and ensures they work better together for high efficiency.

Riverbed SteelHead

- Is the **#1 hybrid** network optimization and application performance solution
- Delivers up to **33x faster** application performance and up to **99%** data reduction
- Allows you to overcome your network speed bumps like latency, congestion, and suboptimal last-mile conditions
- Complements other vendor solutions, like SD-WAN

Our patented solution works in multiple layers, combining:



Application layer optimization, including transfer prediction



Congestion control



TCP stack integration



Data transfer encryption and protection



Scalable Data Referencing (SDR)



Riverbed SteelHead and SD-WAN: Better together

SD-WAN increases efficiencies, while SteelHead:

- Minimizes turns through local connection proxy and look-ahead capabilities
- Reduces the effects of latency and improves the user experience
- Manages ‘chatty’ connections on both server and client communications, limiting the number of conversations that take place over the WAN

Without a doubt, the two technologies work better **together**.




Figure 1: Steelhead handles chatty applications at LAN Speed.

Riverbed SteelHead also works well with industry-standard replication tools, such as:

- EMC Recovery Point
- NetApp SnapMirror
- Data Domain
- Carbonite/DoubleTake
- DAG
- Dell/EMC VMAX
- IBM
- NetApp ONTAP
- MS-Exchange
- AWS
- Azure Cloud

Plus, it can support you across virtual environments like Hyper-V, KVM, and VMware.



How SteelHead can help you meet regulatory requirements

Enterprises across every industry must adhere to regulations in order to stay compliant – and avoid potential prosecution, fines, and costly reputational damage. Healthcare organizations, for example, are required to meet Health Insurance Portability and Accountability Act (HIPAA) laws that dictate how and when records are kept on file.

The Defense Information Systems Agency (DISA) places a tough set of expectations on systems – including a

need for 99.9% availability – while the Payment Card Industry Data Security Standards (PCI-DSS) help protect cardholders and their information.

Similar rules are in place across legal, financial, and military intelligence sectors, to name but a few. There are also uniform protections that all businesses need to employ, like standard contractual clauses (SCCs).

Regulations may dictate a:

Recovery Point Objective (RPO)

The maximum amount of data – as measured by time – that can be lost after disaster, failure, or a comparable event

Recovery Time Objective (RTO)

The maximum acceptable time that an app, device, network, or system can be down after an unexpected incident

Replication and backup time

How often, and how quickly, systems copy and store data

Some, for example, might demand a 24-hour snapshot of all collected data at any given point in time. However, if it takes systems 36 hours to replicate this data, the first batch is outdated by the time the process is finished, and a true picture can never be painted.

This isn't just detrimental to companies' disaster-recovery protocols and decision-making; it also leaves them open to costly non-compliance. The cause? Latency and limitations on maximum throughput.

Uncovering the reality of latency and bandwidth

Latency is driven by the distance data must travel, and compounded when information is distributed over the internet – often over multiple providers. It’s variable, meaning it may change over time. And the higher it is, the worse the user experience.

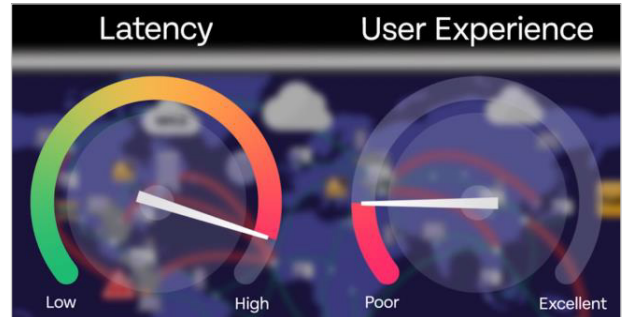


Figure 2: Effects of Latency and User Experience

Traffic traveling from the client to the server and back is called a turn. Turns can occur multiple times per second and enormously impact latency. SteelHead reduces the number of turns through local connection proxy and look ahead capabilities.

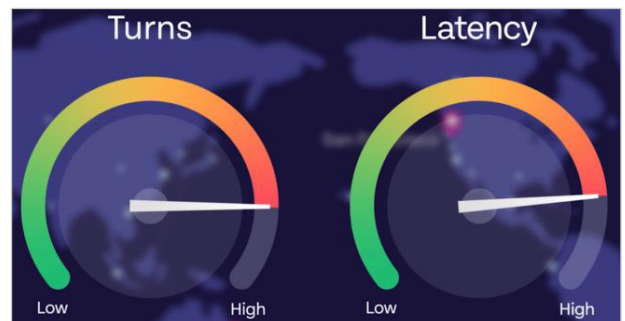


Figure 3: Effects of Turns and Latency



The myth

More bandwidth guarantees higher throughput

The truth

Latency hampers maximum throughput, making increased bandwidth less relevant



In data transmission, network throughput is the amount of data moved successfully from one place to another in a given time period – typically measured in bits per second (bps), as in megabits per second (Mbps). While people often refer to this as ‘available bandwidth’, the

term is misleading. In the example below, you can see that the expectation of using the full 1Gbps ‘available’ in this network is unrealistic due to 150ms of latency. Calculating the actual maximum throughput based on these conditions, we get a figure of 3.495Mbps.

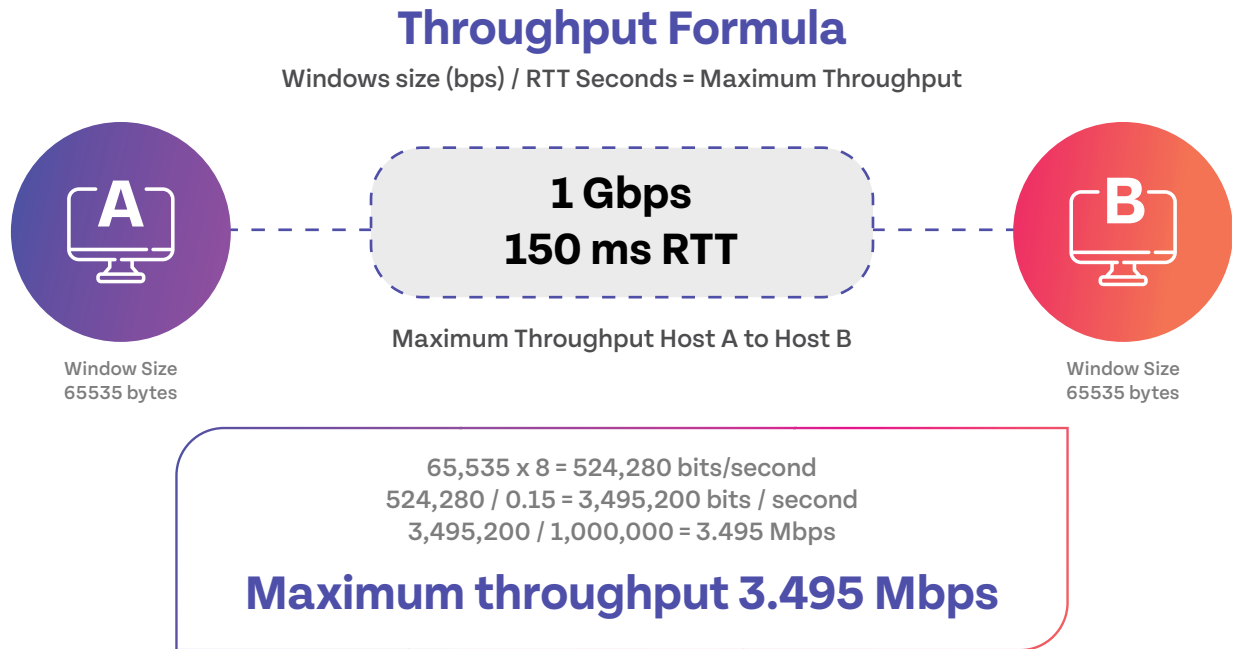


Figure 4: Effects of Latency and maximum throughput

Speed things up to support compliance

One of the most common ways to optimize the speed of a connection is to increase the speed of the link. Links can become overloaded if a device tries to send out too much data; this is called congestion.

HS-TCP

- Designed for high-bandwidth and high-delay networks that have a low rate of packet loss, also known as long fat networks (LFNs)
- Offers many advantages over standard TCP for LFNs, acting aggressively in the presence of packet loss to rapidly maximize bandwidth

SteelHead helps solve this problem by modifying the TCP transmission over the WAN using a unique combination of HS-TCP and MX-TCP.

MX-TCP

- Ideal for dedicated links, or to compensate for poor link quality and high-latency links
- Maximizes TCP by disabling the congestion control algorithm and sending traffic up to a rate you configure, regardless of link conditions
- Fills the pipe to pre-set limits you configure

All this is in contrast to slow-start or regular TCP, which:

- Regulates the amount of data that's sent over a network to stop it becoming congested
- Negotiates the connection between a sender and receiver by defining the amount of data that can be transmitted with each packet
- Slowly increases the amount of data until the network's capacity is reached
- Backs off if congestion is detected

This behavior significantly increases the time it takes to move data.

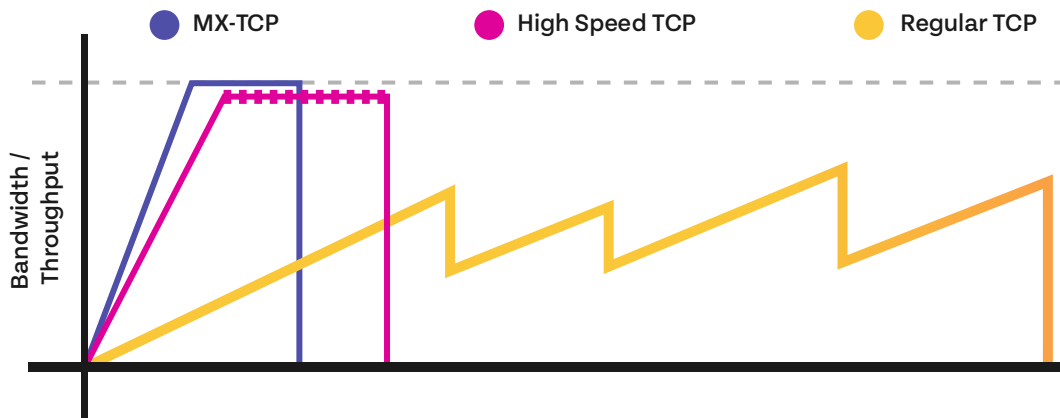


Figure 5: Compare Standard TCP Performance with HS-TCP and MX-TCP

SteelHead's layered, compound approach – which scales linearly – allows it to move data quickly and efficiently, with a sustained data throughput that can consistently exceed up to 90% of available bandwidth. The solution can even be configured to use an exact percentage of your bandwidth to optimize and accelerate data transfer.

So, your data can be backed up faster and smarter than ever before.



Cut bandwidth needs, congestion, and data transfer time with SDR

As well as increasing link speed with the compound approach described above, SteelHead uses scalable data referencing (SDR), a sophisticated deduplication algorithm built to detect repeatable byte patterns in every payload, regardless of origin ‘application’ and whether you’re dealing with files, pictures, documents, or something else altogether. In the majority of enterprise networks, 60-90% data reduction is commonplace.

Even where bandwidth is affordable, Riverbed ensures the least possible data traverses the WAN, reducing congestion. SDR represents the actual data using a 16-byte hash and on warm transfers, which significantly reduces the amount of bandwidth required. References are also scalable, so multiple data sets can be represented with a single hash, reducing the required bandwidth congestion even more. Less congestion and more repeatable patterns mean faster data transfer.

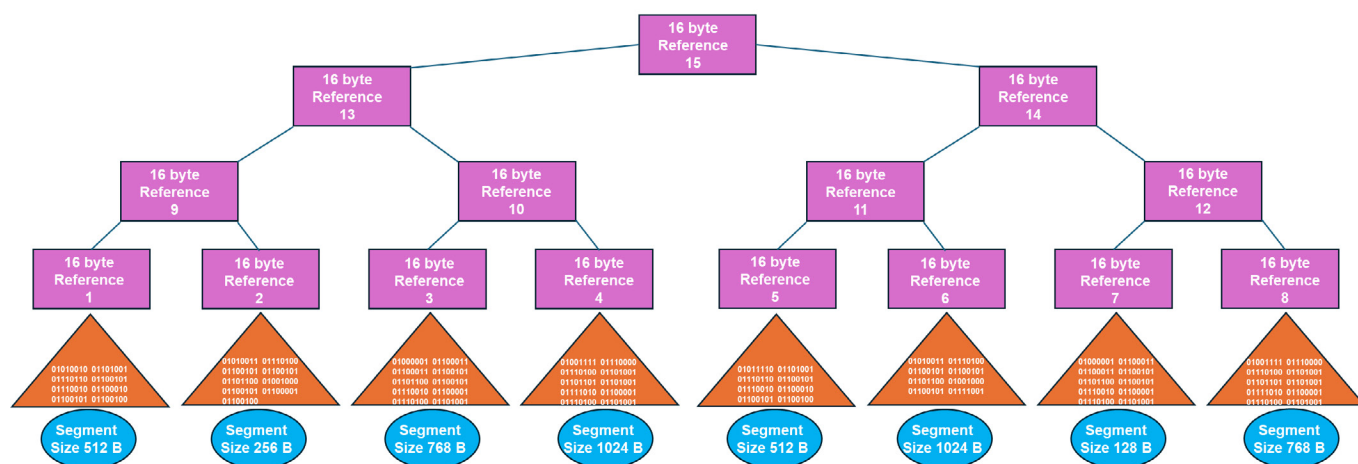


Figure 6: Riverbed Scalable Data Referencing (SDR)



In summary

Riverbed SteelHead quickly and efficiently moves data and scales across multiple clouds and distributed hybrid networks to help organizations achieve efficient RTO and RPO. Customers routinely achieve 90% data reduction and cut transfer time by over 50%. Riverbed SteelHead also works well with industry-standard replication tools.

Acceleration for replication and mirroring not only delivers exceptional ROI, it also supports regulated corporations in achieving compliance goals – a critical need when fines and reputational damage may be at stake.

For more information on these topics and Riverbed solutions, please contact your Riverbed Account Manager or Partner contact.



Riverbed – Empower the Experience

Riverbed is the only company with the collective richness of telemetry from network to app to end user that illuminates and then accelerates every interaction so that users get the flawless digital experience they expect across the entire digital ecosystem. Riverbed offers two industry-leading solution areas – Alluvio by Riverbed, an innovative and differentiated Unified Observability portfolio that unifies data, insights, and actions across IT, so customers can deliver seamless digital experiences; and Riverbed Acceleration, providing fast, agile, secure acceleration of any app over any network to users, whether mobile, remote, or on-prem. Together with our thousands of partners, and market-leading customers across the world, we empower every click, every digital experience. Learn more at riverbed.com.