FEATURE BRIEF: Autodiscovery

Riverbed's autodiscovery capability allows easy deployment of Steelhead appliances without the overhead of setting up tunnels.

Unlike tunnel-based approaches that require explicit identification and configuration of endpoint and tunnel properties, Riverbed's autodiscovery approach provides tremendous flexibility for any type of network – from yesterday's legacy network infrastructure to the latest MPLS-based WAN technologies.

THE WALL STREET JOURNAL
TECHNOLOGY INNOVATION AWARDS
2005

Riverbed Steelhead: Best WAN Accelerator
InfoWorld Technology of the Year

Crucial for Scalability, Riverbed Autodiscovery Lowers Costs and Improves Flexibility Compared to Tunnel-Based Systems

Technologies to improve application performance on the Wide Area Network (WAN) are usually compared in terms of their effect on network traffic or response time – areas in which Riverbed Steelhead appliances excel. Sometimes, however, it is important to consider the costs of initial installation or ongoing management, especially for large networks of devices. Riverbed's autodiscovery capability allows easy deployment of Steelhead appliances without the overhead of setting up tunnels. Unlike tunnel-based approaches that require explicit identification and configuration of endpoint and tunnel properties, Riverbed's autodiscovery approach provides tremendous flexibility for any type of network – from yesterday's legacy network infrastructure to the latest MPLS-based WAN technologies.

Tunnels in Older Network Architectures

Organizations with many locations have often used a hub-and-spoke network architecture for their Wide Area Network (WAN). Each location had a separate dedicated "spoke" link to the "hub" representing the headquarters data center. The largest networks might be composed of a small number of regional hubs, each implementing its own hub-and-spoke network, with a wide variety of higher-bandwidth interconnections between those regional hubs.

Legacy WAN optimization products grew up in this environment, where traffic optimization maps readily onto link optimization. In such systems, a large device or collection of devices is installed into the network at the hub, and smaller devices are installed at each of the various spoke locations. Each spoke device communicates with its hub device via an explicitly-configured *tunnel* for optimized traffic. Each tunnel goes from one site to another across the WAN, and serves as a pipe down which each compressed packet can be sent to the other end.

Newer Networks Are a Problem for Tunnels

However, technologies underpinning WANs have continued to evolve, adding flexibility, functionality and performance that are hard to capture with this older tunnel-based approach. Large networks are often more complex than a simple hub-and-spoke, with multiple hubs. Numerous organizations have deployed any-to-any network "clouds" replacing hub-and-spoke arrangements, and many others are considering such deployments. The new clouds are often based on Multi-Protocol Label Switching (MPLS). MPLS is a next-generation forwarding and switching architecture with important advantages for advanced services and traffic engineering. As MPLS and its relatives become more prevalent in the WAN, it becomes important to understand the architectural implications for tunnel-based systems.

Riverbed Steelhead appliances do not depend on tunnels. While other companies have attempted to tweak, tune, or fix the tunnels they rely on for delivering WAN optimization, Riverbed's approach is entirely different, and offers revolutionary improvements in scalability and ease-of-use.

Any-to-Any Networks: Too Many Tunnels

Consider an organization that has implemented a WAN "cloud," perhaps using MPLS, where any site may communicate with any other site. Such a cloud offers important economy and flexibility in communications because traffic can be routed directly to its destination, rather than having to be transited through the hub site. However, the any-to-any capability of the cloud is a poor match to tunnel-based systems. Tunnel-based systems must establish tunnels prior to optimizing traffic, and the establishment of each tunnel requires some amount of configuration work.

For instance, establishing tunnels typically requires an administrator to identify each and every remote peer tunneling device to each local tunneling device. Then the administrator must configure policies that tell each local tunneling device what traffic to intercept, for every tunnel that will connect to each remote tunneling device. Finally, the same parameters must be configured in the other direction in order to achieve bi-directional optimization. While some tunneling products may have "automatic" tunnel configuration features, nevertheless each of the above configuration parameters must be manually communicated to each tunneling device at least once.

FEATURE BRIEF: Autodiscovery

Riverbed's approach is not only powerful, but fits easily into any network. Riverbed Steelhead appliances using autodiscovery optimize traffic in situations where it can be optimized, and pass traffic through in situations where it can't be optimized – with absolutely no tunnel configuration required.

With Riverbed you can break free of the tunnel vision of other approaches to WAN performance improvement.

Tunnel setup in an any-to-any network therefore is a task that grows in proportion to n^2, where n is the number of sites.[1] Even for systems with automatic-tunnel setup (reducing the cost per individual tunnel) the initial configuration cost for reasonably large deployments can be huge. Worse, adding another site after the initial setup means configuring another n tunnels.

In contrast to the tunnel-based systems, but akin to the network cloud, Steelhead appliances directly support any-to-any communication. The appliances are able to *autodiscover* each other. When using Riverbed's patent-pending autodiscovery capability, traffic between any pair of Steelhead-equipped sites is optimized – *automatically*; there is no need to identify each and every remote Steelhead appliance to each local Steelhead. Similarly, traffic between a site with a Steelhead and a site without a Steelhead is passed through – automatically. The Steelhead appliances are able to deliver this capability not only over MPLS, but also using any technology or mixture of technologies that can support TCP connections: leased lines, frame relay, ATM, etc.

Cost of Full Mesh: Autodiscovery has Zero Cost even as the Number of Sites Grows

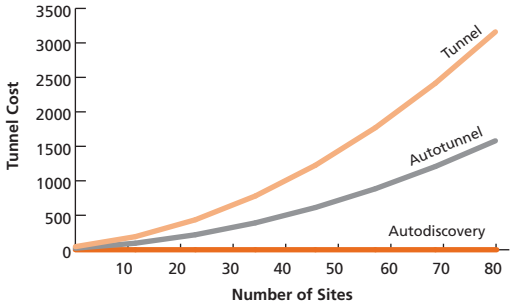

This chart shows the difference between a conventional tunnel-based system, a tunnel-based system with automatic tunnel endpoint discovery ("autotunnel"), and Riverbed's autodiscovery system.

Even assuming the cost of tunnel setup in the "autotunnel" system is only half of the conventional system, for medium-size deployments with 50 sites the costs are enormous compared to Riverbed's zero-tunnel autodiscovery system.

This analysis is conservative because it counts only a single tunnel setup between peer sites. For some kinds of traffic-splitting or fault-tolerant configurations, tunnel-based systems actually require two or more tunnels to be configured between peer sites.

Tunnel-based systems generally have hard limits on how many tunnels they can support. For example, a product line from one vendor supports a maximum of 60 tunnels per unit, even using its largest model. Using that product, it is simply impossible to configure a large any-to-any deployment. Instead, the network designer must choose in advance which 60 peer sites can have optimized traffic. In contrast, Riverbed's autodiscovery system does not have an upper limit on the number of potentially-communicating sites, and can support entirely unpredicted traffic patterns among a large collection of Steelheads.

Not Just WAN Optimization; Not Just WAFS

Of course, Steelhead appliances are about more than avoiding the problems of tunneling. Not only do they deliver application acceleration capabilities in ways that are better than legacy WAN optimization, they also deliver better capabilities. Steelhead appliances are not conventional WAN optimizers; neither are they conventional Wide-Area File Services (WAFS). Instead, Steelheads offer Wide-area Data Services (WDS): broad-spectrum improvement of application performance across the WAN that encompasses the benefits of both WAN optimization and WAFS. By simultaneously addressing issues of bandwidth, TCP latency, and application-protocol latency, WDS systems often produce startling improvements in the end-user experience, especially in difficult network environments that are characterized by low bandwidth and/or high latency and/or high loss.

Riverbed's approach is not only powerful, but fits easily into any network. Riverbed Steelhead appliances using autodiscovery optimize traffic in situations where it can be optimized, and pass traffic through in situations where it can't be optimized – with absolutely no tunnel configuration required. With Riverbed you can break free of the tunnel vision of other approaches to WAN performance improvement.

Riverbed Technology
501 Second Street, Suite 410
San Francisco, CA 94107
Tel: +1 415 247 8800
Fax: +1 415 247 8801
www.riverbed.com

FB-AD110205

[1] The exact cost is $\frac{n^2 - n}{2}$ tunnels connecting n sites.