RIVERBED OPTIMIZATION FOR THE CLOUD

Weaving together the connectivity behind the hybrid cloud

JULY 2011

Changes are afoot in the use and access of IT – hybrid cloud architectures that make the infrastructure more agile by combining the best of private infrastructure with the best of public cloud services are becoming common place. These changes have enormous implications for the mission critical layer of connectivity that brings all of these services together. In this Taneja Group Solution Brief, we'll examine how these changes require a different approach to optimizing the infrastructure, and how one vendor – Riverbed – has extended their deeply innovative optimization technologies. Today, Riverbed is using their technology portfolio to equip businesses with intelligent usability and access optimization that underpins the very performance of the cloud, and makes business use of the hybrid cloud possible.

THE CLOUD VORTEX

It seems as if there's a huge sucking sound coming from the cloud as the industry is pulled into a whirling vortex of new technologies and new ways of thinking that look destined to fundamentally alter IT. The fundamentals of the cloud are simply irresistible to businesses – the promises include simple IT accessibility from anywhere, and integration between IT services everywhere. But at the same time cloud services promise the fruits of the golden chalice in the form of consolidation, better management, and increased efficiency through capabilities such as elasticity (up and down scaling) and "pay as you go" technology use. From this, the idea of public cloud looks inevitably set in motion.

There's little doubt today that the concept of cloud is an irresistible force that is altering the landscape upon which IT systems are built and integrated. The fact is, the cloud is already here in data centers today, as the footprint of IT has rapidly changed and is quickly becoming a mix of services drawn from both public and private clouds. Whether it is by a small degree, or a large degree, the potential of using many of this first generation of Internet-scattered, public cloud services has opportunistically drawn in many users, even within the most stoic and unchangeable of institutions. Meanwhile, the efficiencies, responsiveness, and innovation behind these services are simultaneously changing the way the enterprise builds and deploys their own services with the idea of a private cloud architecture in mind. The idea of a hybrid cloud is this very make up of public and private IT service use, and is the inevitable recipe for IT in the future. The attraction to this hybrid model of IT services is an entirely new level of flexibility and agility that enables IT to pull optimal cost IT services together, and rapidly pivot to new combinations of services as demands change. But it is also planting a tremendous set of challenges before IT organizations. We are living in an age of hybrid clouds, and the hybrid cloud requires an entirely different IT backbone. For the business looking to the future, it is time to get prepared, and it takes a uniquely capable set of technology to do so.

THE LAST MILE IS MORE THAN A MILE

The business network has long been a creature built upon standard practices. Such standard practices have been geared toward optimizing a set of accepted hierarchical connections between users and applications, and optimizing those connections for the most efficiency possible. The hierarchy has often differentiated between ideas of a "core" – the center of high-speed networks that

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557
www.tanejagroup.com

connect key applications and infrastructure – and an "edge" that connects users to IT services. While both are deep engineering exercises, challenges in the edge, often referred to as the last mile, have often plagued the IT practitioner far beyond the challenges that happen in the core. Too often, limited and changing bandwidth, multi-hop latency, limited paths, oversubscription, and unqualified contention still to this day surface in last mile connections that make the most critical aspects of IT – accessibility, usability, and performance - fall apart.

As the business turns to hybrid cloud architectures, connectivity is a different beast to tame, and last mile challenges are popping up in new places. The consequent impacts on unprepared businesses can make next generation IT services fall apart in a world where every mile may be a last mile. Simultaneous combinations of public cloud services like Microsoft Office 365, replication for hosted DR, or many similar Technology-as-a-Service offerings may now fill the previously sufficient corporate Internet pipe with thousands of contending and chatty web applications, making that Internet pipe feel suspiciously like a last mile link. At the same time, a sprawl of virtualized and growing web applications across corporate campuses are filling WAN links with all manner of new collaboration exchanges, simultaneous and redundant session traffic, and multi-point interactions.

Hybrid cloud architectures are creating many more interconnections and types of traffic and any link may suffer from last-mile-like limitations and consequences. The future that holds this inevitable mix of new services scattered across public and private infrastructures makes connectivity a mission critical undertaking. Moreover, the challenges are more difficult than ever, as they are a combination of what may be separate challenges within WAN and remote connections entangled with entirely different challenges within public cloud or Internet connections.

THE CLOUD-OPTIMIZED BUSINESS

In the face of these challenges, there are rapidly emerging and loud demands for solutions and expertise in optimizing business connectivity. It takes the right connectivity to weave the hybrid cloud's diverse set of services together efficiently and responsively, and weave them together in such a way that the business can be intelligent and proactive in adapting to changes in this cloudy world. Such initiatives may become the dominant activity for IT departments as the transition to this next generation of IT architecture takes place.

Cloud-optimized businesses will require more than fast pipes. They will also require an integrated combination of technologies that are designed to simultaneously address private WAN optimization and public cloud service acceleration and the complex intersection of those two technologies beneath major sets of strategic IT services. We see three fundamental pillars that must be addressed in bringing these capabilities together for the right infrastructure foundation beneath the hybrid cloud:

Connection. The cloud is about connection, and more connection than ever before. Some cloud traffic will inevitably be more important than other traffic, and this will constantly change. Amidst this whirlwind, connections will require optimization to work across limited physical links. But optimization will have to separate the wheat from the chaff, reach deeply into the protocols that make up a multitude of connections crossing every wire, and constantly react to changing services to deliver a business class cloud. The connection is fundamental in the age of the cloud, but it is more complex and demands deeper optimization than ever before.

Insight. Connection in the age of cloud is more dynamic than ever, and in the midst of dynamic connections the business requires insight and a platform for proactive application-aware performance management. Irrespective of how well optimization may adapt to changing cloud uses, businesses will need to interact with connectivity more than ever before. Where storage systems and servers used to be the foundational layer of IT, with the shift to cloud solutions, connectivity will now serve as that foundation. As a critical and precious resource, connectivity will require equal if not better insight, monitoring, alerting, and management.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

Content. While the connection is fundamental to use of the cloud, interaction with the cloud is more content-centric than ever before. Much like we have done with ever faster processors and ever cheaper memory, inefficient connectivity solutions may deal with a deluge of content by throwing more connections and bandwidth at it. But connections are not ruled by Moore's law, and involve more physical and market limits. Optimally delivering cloud services requires a foundation for efficiently bridging the gap between where content originates, and where content is used, with application-level awareness. Anything short of application-awareness runs the hazard of failing to optimize chatty, duplicate data exchanges, and will fail to localize the experience of interacting with this content-centric, but distantly removed cloud. Moreover, failure to fully optimize in the best way possible may mean the difference between realizing the cost savings promises of the cloud, or simply converting expensive IT systems into the cost of cloud service data transfer.

The hybrid cloud is a tremendous blender of digital interaction, and it takes distributed, intelligent, and automated optimization to contour the resulting flurry of digital bits into business class services. These three pillars woven together with a well integrated, distributed architecture can place the right data in the right place at the right time, and overcome the challenges imposed by the cloud. But weaving these three pillars together in an integrated solution set requires a distinctly different and more comprehensive approach than traditionally applied to private WAN business connectivity.

THE RIVERBED OPTIMIZED HYBRID CLOUD

In this age of the cloud, we are constantly casting about the market looking for vendors poised to deliver ideal integrations of cloud optimization technologies, and Riverbed has long proven that they are uniquely equipped to deliver just such a solution set. As a pioneering innovator in the WAN optimization market, Riverbed has long reached beyond fast pipes to deliver optimized data interaction that can make the most remote set of data or applications perform like a local application. Riverbed's optimization technology goes far beyond pure compression, and makes the most of every opportunity to take inefficiency out of application and data interactions, and such technology can transform both public and private cloud services into efficient business solutions.

Cloud Connection

With a shift to applications and data from anywhere, the cloud requires connection optimization that reaches into applications and data to deliver access and usability that matches a local experience from anywhere. Today, Riverbed's connectivity optimization for the cloud revolves around two technologies, Steelhead products and Whitewater. Let's take a look.

THE MULTI-DIMENSIONAL OPTIMIZING STEELHEAD FAMILY

Riverbed's Steelhead products are delivered as hardware appliances, cloud and virtual appliances, and client software that optimize data interactions across the corporate network and the public Internet. In one dimension, the power in Steelhead products comes from a layered optimization engine designed to optimize every transmitted bit. Steelhead optimization products deduplicate data transmitted over the wire, tackle inefficiencies in network transport interactions, mediate and optimize application interactions to squeeze out redundancy and chattiness, and manage the full protocol stack through a QoS engine that auto-adapts and self-tunes optimization under changing network conditions.

But while the optimization from Steelhead products are often touted as accelerating data and application access by factors of 10x or more, another dimension of optimization is found in the varied flavors in which Steelhead products can be deployed – a variety that lets optimization reach anywhere. Steelhead products can be deployed in forms ranging from enterprise data center scale appliances to remote office appliances, to virtual and public cloud versions, and even software clients that can optimize performance for the remotest of remote users. The Steelhead portfolio delivers in-

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

depth optimization, and in-breadth touch that can stretch across any business network – from the laptop to the virtual infrastructure of public cloud providers such as VMware and Amazon.

WHITEWATER

Riverbed has extended their data optimization capabilities to cloud data storage as well. Whitewater is a deduplicating cloud storage backup appliance that sits in the customer data center and seamlessly connects a business to low cost storage being served up in the cloud today by Amazon S3, Nirvanix, AT&T Synaptic or other EMC Atmos-based service providers (including on premise private clouds). Moving beyond on-premises deduplication, Whitewater pairs together low cost cloud storage and Riverbed's proven deduplication technology to further reduce cloud storage utilization and cost, while delivering local-storage-like responsiveness, connection-loss tolerance, and practically unlimited scalability within what is today a 3U physical footprint. Whitewater is deployed as a deduplicating disk-to-disk backup target behind any of the leading backup solutions – including Symantec's BackupExec and NetBackup, CommVault's Simpana, or EMC's Networker. As backup data is written to Whitewater, Whitewater connects with the high redundancy and nearly infinite storage capacity available from public cloud storage services to scale backup storage to any capacity, while simultaneously deduplicating transmitted *and* stored data, so that low cost cloud storage becomes *extremely* low cost cloud storage. Meanwhile, the stateless Whitewater appliance caches local data to make backup storage tolerant of any cloud connection outages and make the cloud perform like a local service. This statelessness then turns a single Whitewater backup device into an all in one disaster recovery solution as well – Whitewater data in the cloud can be recovered with any other Whitewater appliance. In turn, Whitewater has become a cloud storage gateway that is equal parts about optimizing backup performance and capacity utilization, and about reducing data lifecycle costs with an all-in-one solution for long term data preservation.

Cloud Insight with the Cascade Suite

Amidst the onslaught of new connections and digital interactions that make up the hybrid cloud, Riverbed has steadily expanded their expertise in optimization into expertise in insight and management. Riverbed has assembled a full-blown application-aware performance management suite they call Cascade. Cascade is a family of Riverbed software and appliances (Shark, Pilot, Sensor, Profile, and Gateway) as well as Riverbed Steelhead technology that can peer into the application-level interactions taking place in the infrastructure, and help organizations use Steelhead technology to reshape traffic and mitigate on-going performance challenges. When it comes to the hybrid cloud, these products can help organizations identify application interactions and dependencies as a precursor to migrating applications to cloud architectures, or transferring application components to the public cloud. Within the hybrid cloud infrastructure, the Cascade suite rolls up the traffic patterns captured by Riverbed's portfolio of network flow, application, and network analysis devices – Shark, Pilot, Sensor, and even Steelhead family technology – and presents current network conditions in dashboard fashion so IT organizations can identify whether network characteristics are operating in accordance to Cascade-defined service level metrics, and whether parts of the infrastructure may be trending toward out-of-norm conditions. In response to out-of-norm conditions, the Cascade suite can alert IT shops to proactively take action before services become unresponsive and help them drill down into interactions for troubleshooting. Moreover, IT departments equipped with Cascade can proactively identify opportunities to rearchitect IT services in the future and turn to new cloud services, with the right expectations, as they emerge. With Cascade, Riverbed is equipping organizations to constantly reshape their connectivity and approach to IT service delivery. This is especially important behind cloud-oriented businesses.

Content Optimized with the Steelhead Portfolio

Finally, while optimized transmission and interaction solve many challenges in the era of the cloud, transmission alone shouldn't be the only weapon in the arsenal of tools deployed to make the cloud

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

business-ready. The cloud is content-centric, and to keep up with the onslaught of redundant content that is flooding pipes, vendors will need to innovate around how to bridge local access to remote content without chewing up valuable and expensive network resources. To speed access to cloud content, Riverbed has long been pioneering data reconstitution from deduplicated data elements previously transmitted across the wire. The payoff has been performance efficiency that enhances usability and access no matter the link. Moreover, with Cloud and Virtual Steelhead appliances and endpoint clients in tow, Riverbed has been able to deliver that performance from any data center to any endpoint and cover all corporate users utilizing corporate applications – the ultimate tool for delivering services from the private cloud.

But more recently, Riverbed has announced their intent to extend their optimization reach, and connect it to Internet distributed content through a partnership with Akamai. Akamai's technology has long been behind the caching and delivery of content, including the content within cloud-based Software-as-a-Service (SaaS) applications, with an internet-edge architecture made up of proprietary algorithms running across more than 90,000 caching servers scattered around the Internet. Today, Akamai's solutions include SaaS-specific, across the Internet technologies like dynamic service mapping, dynamic routing, and connection optimization that make services from the distributed public cloud more responsive to the business customer. With a Riverbed and Akamai partnership, the future looks bright for bridging Internet content caches into the corporate data center, to massively reduce bandwidth, session traffic, and needlessly retransmitted data. Meanwhile, the power of Riverbed's optimization reach will be extended well outside the data center walls, and help businesses gain control of how their traffic is optimized across every Internet link to nearly any service and to nearly any endpoint.

TANEJA GROUP OPINION

Shifting to the next generation of IT services – made up of a hybrid mix of IT solutions drawn from the private infrastructure and the public infrastructure – will require nothing short of a wholesale reinvention in how companies peer into and manage their infrastructures, and in no domain will this have deeper impact than in the network. Where the network in the past has been architected around center to edge, bi-directional, and largely unchanging traffic patterns, the network today is awash in an anywhere to anywhere mesh of dynamically changing traffic that will only become more dynamic in the future.

In the midst of this complexity, network optimization innovators will even more clearly standout, and the laggards with point solutions will fall away. It is clear that Riverbed has constructed a broad, deep, and innovative technology portfolio designed to not only optimize any traffic, but to deliver the right performance, in the right place, at the right time through an integrated assemblage of multi-layer optimization, network performance management and content acceleration technology. In a way, Riverbed was traveling this road long before the idea of cloud came about – Riverbed after all cut their teeth on optimizing the remote and branch office across the WAN. In a very similar way, the WAN has long been about enterprise consolidation moving resources further from the branch and end user. Strategic approaches to cloud architectures will require the right foundation, and as of today Riverbed delivers the only portfolio we see that reaches deeply and broadly enough to provide that foundation.