



POWERING THE HYBRID ENTERPRISE

TO BETTER SERVE BUSINESS DEMANDS FOR INFORMATION EVERYWHERE, ENTERPRISES MUST DEVELOP NEW STRATEGIES FOR OPTIMIZING MULTIPLE KINDS OF NETWORKS.

The panes of the glass house have long been shattered, leaving different kinds of pains in their place. Computing is no longer done in a single location. Users, applications, and data exist in more places than ever before, creating an unprecedented challenge for IT.

Numerous Forrester Research reports make the evidence clear. Its surveys show that 62 percent of employees work in multiple locations,¹ that 51 percent of IT professionals say application complexity is their primary challenge,² and that 52 percent of enterprises have more than half of their corporate data outside the data center.³ To deal with these challenges, IT must determine how to overcome myriad technical constraints and allow anytime, anywhere computing.

In addition, by accelerating delivery of data and creating optimal user experiences, IT can help enterprises increase employee productivity. IT can also ensure that the data employees need is accessible in systems with the minimum amount of latency based on where they work.

How can IT achieve the flexibility and agility it needs to offer multiple types of applications in multiple locations? Just as enterprises take advantage of applications hosted in multiple locations (i.e., public clouds, private clouds, or SaaS applications), they take advantage of different networks for different reasons (i.e., performance, backup, remote access). These could be MPLS links, ATM links, traditional Internet connections, or VPN lines. Enterprises, in effect, are already building hybrid networks, and those combined with the elements of hybrid clouds create the hybrid enterprise.



The result is a wide-area network (WAN) strategy designed for optimal performance while minimizing recurring cost. Just as hybrid clouds simplify the transfer of data, so too do hybrid networks. But hybrid networks do more. They deliver a crucial element to IT, allowing it to shift and manage the transfer of data to the most appropriate network based on need. No matter what the application, no matter what the network, IT can direct data to the most high-performance or cost-efficient link.

The reward for this evolution is not only an unprecedented level of dynamism for the network and agility for the enterprise, but the ability for IT to apply the advantages of on-premises applications — SLAs, bandwidth, acceleration — to cloud applications.

/// BUILDING THE HYBRID ENTERPRISE

Enterprises already take advantage of multiple networks. In its efforts to match application need to application performance, IT maintains multiple networks. These may be private networks for critical business services, comprising point-to-point links, ATM frame relay networks, or networks using the MPLS protocol. All these provide a variety of benefits: performance, security, and, of course, control.

Enterprises may also employ Internet connections to cloud services and other online resources. These involve a trade-off: They are less expensive than private networks, but they are also less reliable.

While these diverse WAN links provide flexibility, they can also cause problems. IT lacks visibility into them, and thus can't control them as efficiently as it would like. In fact, according to research firm Gartner, 70 percent of IT professionals say they can't easily diagnose outages, application issues, link failures, and other network disruptions. IT might be willing to maintain the status quo regarding its management and monitoring of these links, but unfortunately, new technologies are changing and challenging IT strategies for dealing with multiple networks.

IT needs to adjust to a new reality when it comes to applications running over the network. While it may have been accustomed to managing traditional business applications run over private links, half of network applications today are running on public clouds, mobile devices, or websites. What's more, Gartner notes that 80 percent of traffic running on WANs is in large part because of mobile and SaaS applications. As a result, IT must rethink its definition of a WAN.

Shifting business demands make hewing to the status quo seem impossible as well. It's not just the increasing demands of a highly mobile workforce. As if serving those mobile needs weren't challenging enough, enterprises are demanding more online collaboration and video conferencing. These demands further stress networks, as end users expect to exchange documents, graphics, and other rich media in real time.

/// FOUR THINGS IT MUST HAVE TO MAXIMIZE HYBRID NETWORKS

The path is clear. To accommodate new business demands, IT must institute changes in the way it manages and monitors its networks. Indeed, Gartner's 2013 report "Is MPLS Dead?" noted: "Network architects should revise WAN architectures to improve performance for external cloud applications and resources. In most cases, hybrid WAN architectures will provide the best blend of performance and availability."

What kinds of shifts in its capabilities should IT be thinking about? They coalesce around four categories.

END-USER EXPERIENCE. IT needs to ensure instant, reliable performance for end users directly accessing Internet and other cloud resources.

HIGHER COST EFFICIENCY. Spending less and getting more may seem counterintuitive, but it can be done. Traditionally, IT teams have typically overprovisioned networks to ensure that bandwidth is available. But there are ways to ensure availability

No matter what the application, no matter what the network, IT can direct data to the most high-performance or cost-efficient link.

by using bandwidth the organization is already paying for, and to avoid paying for standby backup links when they're not necessary. For example, for some requirements, enterprises can significantly reduce monthly recurring WAN costs by running existing MPLS links over the public Internet.

INCREASED RESILIENCY. To accommodate the first two, though, IT must increase resiliency. It needs to better detect outages, so that it can shift applications when necessary on a more dynamic basis. It must be able to ensure availability and reliability of priority applications, maintaining uptime and improving time to remediation.

INCREASED FLEXIBILITY. IT also needs the flexibility to limit certain applications when outages occur, so that remaining circuits are open to mission-critical applications. For instance, compliance regulations dictate that PCI data must run on qualified networks, so it would be given priority.

Overall, to achieve these goals, IT needs higher levels of *visibility*, *control*, and *responsiveness* to deliver optimum use of network resources to the business.

To improve its success rate, IT needs even more granular insight into the network. This insight needs to be more proactive than reactive, with an understanding not only of what's happening but what needs to happen. This requires developing richer knowledge about application and network traffic; in fact, it means integrating application- and network-awareness in order to create intelligent policy configuration and troubleshooting. It also involves developing a capacity and traffic engineering plan for data and a method of measuring the effectiveness of network links once they've been deployed.

In addition, IT needs a higher level of control when it comes to configuration and deployment. This means aligning applications with the optimum network paths available. It means setting and automating policies for prioritization, network paths, resources, and failover. Furthermore, it means backing up these

HOW RIVERBED HELPS CREATE THE HYBRID NETWORK



Visibility



Control



Optimization



Simplicity

Riverbed offers a wide variety of highly integrated technologies that allow IT to optimize, monitor, and control business-critical applications to ensure that end users get the performance they need, while reducing the risk of application downtime. Its tools also deliver increased visibility and control for businesses with a high number of branch operations, but without the expense of having to add dedicated hardware.

As a result, IT derives multiple benefits. It gains greater visibility into applications and control over managing the end-user experience of the applications on which the business runs. It can allocate application delivery over the most appropriate network, moving applications with higher priorities onto faster networks while shifting lower-priority applications — such as backups and updates — to less-expensive networks. IT can now centrally manage and automate QoS policies to ensure it isn't under- or overprovisioning bandwidth, as well as monitor networks in real time to ensure availability.

For more on how Riverbed optimization tools can help enterprises, visit www.riverbed.com.

policies for application performance with QoS and service-level agreements. Finally, greater control means understanding application performance regardless of whether the applications run on-premises or in the cloud.

The upshot for IT is the ability to be more responsive to the business. It should be able to deliver optimal access to applications and data, regardless of the location of the end user or the service. It should be able to maximize throughput for data tied to the most stringent SLAs. It should be able to increase end-user performance with application optimization. And, perhaps most challenging, it should be able to provide scale and performance for both physical and virtual platforms.

All of these issues affect IT's ability to accommodate changing business conditions. IT must be able to centrally and intelligently map the network so that it is not only aware of the applications (and in the case of many enterprises today, this can mean thousands of applications that live both on-premises and in the cloud), but of the applications' relation to business priorities. And it must be able to apply more automation when making these changes, eliminating delays that stem from manual configuration and human error.

/// GETTING THE MOST FROM THE HYBRID ENTERPRISE

With the strong foundation of a hybrid network in place — and especially the new levels of visibility, control, and responsiveness they bring — organizations can then move toward becoming a true hybrid enterprise. This encompasses three key areas.

ANALYTICS. With visibility into the network comes the ability to collect, synthesize, and process application performance data. That analysis, in turn, provides IT better support for the planning and simulation of new network links.

DIAGNOSTICS. With the ability to see how applications move across mixed topologies and hybrid architectures, IT can better inform cross-functional and DevOps workgroups. This delivers insight and understanding into how the development of applications and the deployment of applications tie together.

TRANSFORMATION. With greater visibility, enterprises can transform the way they deliver applications and data — with more acceleration, elasticity, and scalability across centralized points of automation and control. They can thereby support all manner of hybrid architectures, whether they encompass SaaS, cloud, or Web applications, and be more agile as business requirements change.

Enterprises are already deploying the initial stages of hybrid networks, and are well on their way to creating the hybrid enterprise. They're deriving

multiple benefits in terms of speed, performance, and availability.

Manufacturer Michelin, for instance, faced complexities in application troubleshooting across its global networks. With state-of-the-art tools, it gained end-to-end application and infrastructure visibility and was able to better understand how applications would work on the network during development. The result: a 30 percent reduction in the time it takes to roll out applications.

Energy management company Schneider Electric faced a lack of control over the Internet delivery of cloud applications. With the proper tools, it's getting better performance, reducing latency to give LAN-like performance on WAN-based networks, and better control of cloud applications.

Engineering conglomerate Siemens had to deal with costly IT infrastructure at a remote manufacturing site. By improving its network capabilities, it was able to improve its resource utilization. It reduced its IT footprint at the remote site from five servers to one.

Finally, financial services firm Spot Trading faced retention requirements that required significant off-site storage, but the amount increased so quickly that its IT team found that it took days to back up and days to restore. By deploying tools that optimized its cloud connection, Spot Trading avoided the \$500,000 cost of upgrading to a larger SAN on-site and reduced its recovery time of files from days to minutes.

These companies exemplify how enterprises are demanding more from their networks than ever before. They need flexibility and simplicity for their current activities, but more importantly, they need a way to maintain flexibility and simplicity going forward, when business demands on applications, the network, and infrastructure are only going to increase. Ultimately, enterprises must focus more assiduously on how they can optimize all their new hybrid WAN connections — from remote sites to the cloud and beyond.

IT must be able to centrally and intelligently map the network so that it is not only aware of the applications (and in the case of many enterprises today, this can mean thousands of applications that live both on-premises and in the cloud), but of the applications' relation to business priorities.

¹ Forrester Research, "Provisioning a Flexible Workplace Pushes the Boundaries of IT and Drives a Search for Outside Help," June 2012

² Forrester Research, "Think You've Mastered Application Performance? Think Again," July 2013

³ Forrester Research, "Successfully Consolidating Branch-Office Infrastructure in the Face of More Users, Services, and Devices," October 2011