

How Poor App Performance Impacts the Enterprise in the Cloud Age

How to address It

Poor application performance is nothing new. But in the era of cloud computing, as-a-service IT delivery models and 24/7 digital businesses, its potential impact is greater than ever. This paper looks at the causes of poor application performance, the business implications and some ideas on how to address the problem before it throttles organizational effectiveness.

No one needs a lecture on the growing and integral nature of technology to all organizations. But not everyone may realize just how large the application tsunami has become. For instance, research indicates that the number of software applications deployed by large enterprises jumped 68% over a recent four-year period, with many enterprises now utilizing well over 200 applications.¹ Even small organizations are now juggling nearly 75 applications in their software portfolio.

And it's not just about application volume. The data generated by those applications is increasingly viewed as mission-critical, running everything from kidney infusion pumps and electronic braking systems to split-second financial transactions valued well into the billions of dollars.

Against those twin developments—application proliferation and data primacy as a strategic asset—a major paradigm has emerged: the negative impact of poor application performance. Poor application performance has multiple causes, and the business impact of poor application performance can range from frustrated users and customers to loss of revenue and diminished brand reputation.

The problem, unfortunately, becomes more pronounced as more applications and more vital data are used by enterprises. This is further exacerbated by the increased prevalence of performance-limiting trends such as cloud computing, web-based applications and global supply chains being added to the mix.

While organizations are looking at numerous options for dealing with unacceptable application performance, enterprise decision-makers need to do their homework to understand the causes and impact of application degradation. And they must understand that not all solutions really get to the root of the problem driving application performance—network and transaction latency.

Networking and IT decision-makers are looking to address these problems in a more efficient and scalable manner than they have been able to do in the past. This paper also will cover the importance of modernizing their perspectives and deployments of WAN optimization, in order to go beyond simply addressing bandwidth and type of network. The proliferation of applications—especially cloud-borne applications—means that a new view of WAN optimization is essential to properly address the issue of subpar application performance.

Networking and IT decision-makers are looking to address these problems in a more efficient and scalable manner than they have been able to do in the past.

Why application performance matters more than ever

Because applications are only as important as the data and insights they produce, there is more pressure than ever on organizations to ensure that the increasing number of applications—on premises, in the cloud or both—available to users is actually delivering value. Otherwise, application proliferation dramatically increases cost, but it also significantly impacts responsiveness and data availability.

There are a number of critical reasons why ensuring suitable application performance is more important than ever. First, there are the “experience” concerns, primarily user and customer experience. We’ve all felt the rising frustration levels when we can’t access our preferred applications and, by extension, use the important data we need to make smart decisions. That user/customer

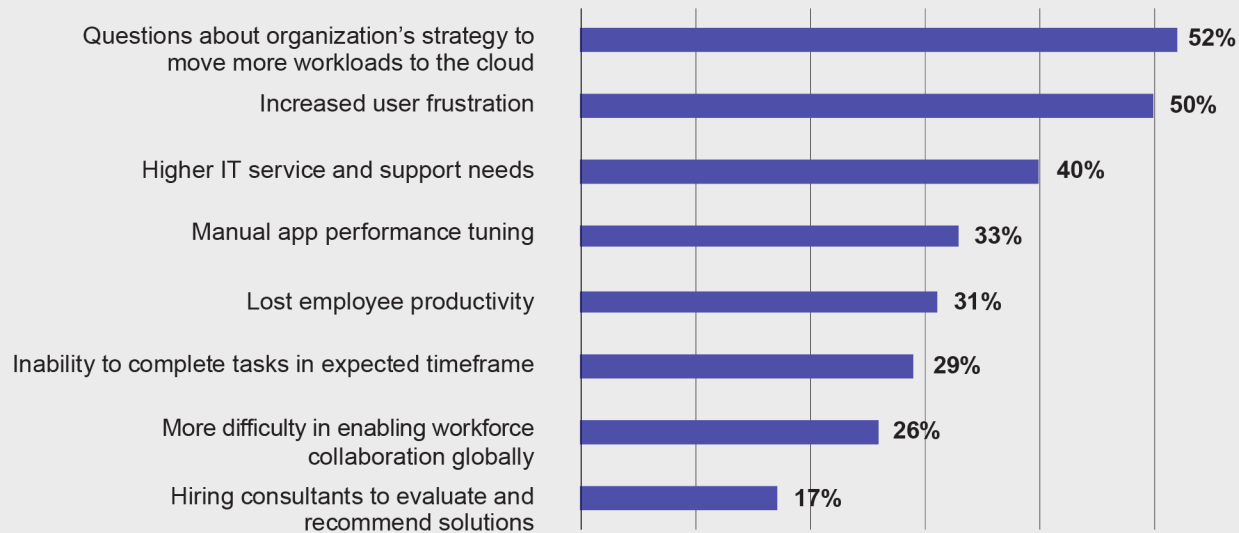
experience must meet or exceed expectations—no easy task when you consider that application access must be provided anywhere, at any time. This is not your legacy computing model, where the vast majority of users accessed applications on a single machine, from a single location and typically during “normal” business hours.

An important extension of the user/customer experience rears its head when performance lags badly, causing consumers to abandon online shopping carts and employees to drop efforts to collaborate with colleagues because latency drags meetings, file sharing and conversations to a dead halt. Not only do problems like that cause problems in the areas of customer retention and lost employee productivity, but they also impact the organization’s brand perception if it is seen as a technology laggard or nonresponsive.

Additionally, service and support—critical to how customers interface with a brand—can suffer tremendously by poor application performance. Think of the mindset of a buyer calling the customer service desk over a technical or billing issue or an employee trying to get a help desk ticket resolved so they can quickly get back to work. Application performance—whether caused by latency, network bandwidth or undue complexity—has the potential to make or break a service call.

Finally, one of the key areas that is impacted by subpar application performance is collaboration, an increasingly strategic and vital part of all organizations’ operating procedures. With workforces increasingly global and mobile in nature, access to the same applications and data sets is essential to promote collaboration throughout the organization. Think how frustrated you get when participating in a virtual chat room or a teleconference and you don’t have access to the same data at the same time as other participants. Your interactive session has now become analog, productivity suffers, frustration grows and employees lament their organization’s inability to fix the problem. And think about the frustration, lost productivity and missed business opportunity caused by a global R&D team being unable to overcome crippling latency while sharing new design specifications for a strategic product launch.

Which negative business impacts have you experienced as a result of Office365 application-performance challenges?



In research recently conducted with IT and networking decision-makers who are registered members of TechTarget's universe of technical buyers, respondents said Office 365-related application performance issues have the potential to undermine their organization's "cloud-first" strategies. This data points to a heightened awareness to, and concern over, the impact of unacceptable application performance.

While some organizations may have grown accustomed to cloud collaboration as a reason to settle for less-than-ideal performance, whether through complacency or an inability to properly address the problem, clearly that is starting to change. Organizations can no longer accept less-than-ideal performance as an inhibitor to successful digital transformation, especially given the investments in money, time and personnel that are being made.

Causes of poor application performance

For years, it has been natural for organizations facing application performance problems to view bandwidth limitations and/or complexity as the key culprit. And there's no doubt that bandwidth challenges do exist and

do limit organizations' ability to provide suitable application performance levels. While this is undoubtedly influenced by application proliferation and the resultant explosion in the volume, variety and velocity of data, network bandwidth cannot be targeted as the sole—or even primary—source for poor application performance.

It is also true that infrastructure complexity adds to the troubles. Today, all enterprises have a mix of infrastructure types in both physical and virtual environments, and the longtime problem of server sprawl now has given way to VM sprawl, as users quickly and cost efficiently stand up their own virtual machines. As enterprise networks increasingly are required to support hybrid cloud environments, managing those networks becomes more complex, while also opening up potential blind spots to the IT organization. The resultant complexity breeds inefficiency, which in turn impacts application performance.

Another key contributor is the rise of tech-savvy (but often impatient) employees using their own applications and accessing them through their own devices over the enterprise LANs and WANs. This phenomenon, called

shadow IT, often is an extension of corporate-sanctioned BYOD/BYOA policies, but it also can result in more complexity and more applications vying for network resources when users ask for access and data.

Finally, however, it is important to keep in mind the primacy of unacceptable latency as a performance inhibitor. Organizations can get everything else right—suitable bandwidth, reduced complexity, more control over rogue IT activities—but it is escalating latency that brings application performance and employee productivity to its knees. In an increasing number of real-time applications where delays of milliseconds are vital—think of high-speed electronic trading or electronic monitoring of healthcare problems—latency is a drag ... literally.

And in an increasingly real-time-only world fueled by pervasive mobility, cloud computing, 24/7 e-commerce, global collaboration, SaaS and cloud computing, organizations need help in eliminating latency that slows application performance. A harsh reality is that predicting latency is next to impossible, given the variety of locations and spans of distances where application access is critical. It doesn't matter if you're trying to access Office 365 or another mission-critical app in a


headquarters facility, home office, customer site, coffee shop, airplane or anywhere in the world; latency is likely to be the biggest threat to application performance.

Case in point: Office 365 application performance

Office 365 is a huge success story; it's one of Microsoft's most successful software products ever, and many of the core applications within the suite (such as Office apps, SharePoint, OneDrive and Teams) are now considered to be strategic by enterprises. Another element of this increasingly strategic role Office 365 now plays within most enterprises is its valued position as the "on-ramp to the cloud," particularly in light of the increased move toward edge computing.

This has created an important paradox, however, in that users now expect Office 365 applications to respond instantaneously, so any performance lag is noticeable and annoying and impacts business operations. In many ways, the pervasive use of Office 365 and its increasingly strategic utilization throughout the enterprise has forced users and decision-makers in IT and networking to find new approaches to deal with subpar application performance.

Which have been your greatest concerns since adopting Office365



For instance, TechTarget research points out that Office 365 users are, in fact, taking notice of unacceptable application performance and believe that failure to address those performance problems could have a deleterious impact on their organization's plan to move more applications, workloads and data to the cloud.

The above graphic highlights numerous performance issues users have experienced with Office 365 applications.

While it should not surprise anyone that security remains a critical issue for Office 365 users (and, in fact, for most SaaS and cloud-based applications), performance issues loom large in the minds of Office 365 users surveyed. In fact, most respondents selected at least two different performance-related issues as a concern for using Office 365.

Latency, in particular, remains a big sticking point for Office 365 users. Among the survey respondents, 60% said it was either "very important" or "essential" for their organization to adopt tools and techniques to address latency issues with Office 365. And 94% of those respondents either agreed or strongly agreed with the need to achieve the shortest, secure path via the internet and to also have an optimized connection to Office 365 applications for consistently good performance at all times.

The increasing popularity of SaaS and cloud-borne applications is a stark reminder of just how much the world has changed when it comes to how applications are consumed.

The benefits of modern WAN optimization

In their heightened efforts to improve application performance and provide faster, more reliable access to data in real time, many organizations have tried multiple approaches to address the problem. These include buying more network bandwidth, installing automated network management tools, finding new ways to more efficiently direct data packets over the network to and from the cloud and WAN optimization/application acceleration tools.

Recently, organizations have extended their adoption of software-defined solutions with SD-WAN, an important approach that simplifies network complexity, reduces costs and eases management. But many buyers of SD-WAN solutions adopted the technology because they mistakenly felt legacy WAN optimization tools weren't enough to deal with application performance limitations. IT leaders and networking buyers, however, should be aware that few of today's SD-WAN solutions are actually designed for the new reality of application performance. For enterprises looking to address application performance, it's important to view modern WAN optimization as a complementary, yet integral piece of a performance-driven enterprise. In fact, there are WAN optimization solutions available today that address the full gamut of application environments, including on-premises, private cloud, hybrid cloud and, of course, web-based SaaS applications, while also addressing modern consumption models that span client-based mobile users for application performance anywhere in the world.

The increasing popularity of SaaS and cloud-borne applications is a stark reminder of just how much the world has changed when it comes to how applications are consumed. This makes network performance unpredictable, especially as workforces become more mobile, dispersed and virtual.

And even organizations that have embraced SD-WAN now realize that the technology intended to modernize network management and provisioning does not replace but rather must coexist with WAN optimization.

Research from Frost & Sullivan, for instance, points out that 67% of U.S. IT decision-makers still consider WAN optimization essential and would like it to be integrated with SD-WAN environments.²

Data from the TechTarget research provides a clear indication of respondents' desire to have a more modernized implementation WAN optimization tools in order to more fully address application performance challenges.

Bottom line: WAN optimization and application acceleration tools can help address these issues, but organizations must take an updated, modernized look at those tools in order to ensure that they are properly addressing the core causes of subpar application performance. SD-WAN can help to a certain extent, but the optimal approach is one that combines SD-WAN with modernized WAN optimization approaches to overcome the latency challenge that most acutely impacts application performance.

Riverbed: A leader in application performance for the modern enterprise

A longtime leader in WAN optimization and application acceleration, Riverbed Technology, Inc. is uniquely positioned to help organizations address the increasingly difficult task of improving application performance for global enterprises in a mobile, distributed, virtual, cloud-first, real-time data operating model.

Riverbed® not only offers organizations both a broad and deep range of easy-to-implement services and solutions optimized for today's application performance challenges, but it also has staked out a leadership position in the modernized view of WAN optimization. This makes Riverbed's solutions well positioned for organizations that depend upon SaaS applications (such as Office 365) to meet daily business requirements, any version of cloud computing, and those that are interested in application performance working in symmetry with enterprise-class SD-WAN technology. Riverbed also offers a full suite of network performance management tools designed to provide end-to-end visibility into network and application performance.

Riverbed's solutions help drive performance anywhere an organization's applications exist—on premises, in the cloud or both—and anywhere users are consuming them. A big plus for organizations partnering with Riverbed is the company's many years of experience in the way applications and data behave over networks, and they understand that not all WAN optimization tools are created equal. In today's modern enterprise, it's essential to incorporate tools that will optimize and accelerate performance of the network and applications for a fast, reliable and consistent user experience. This undoubtedly enables an improved competitive posture for enterprises, while putting them in a position to achieve business outcomes more efficiently, reliably and securely.

Conclusion

Application performance is a critical issue that IT and networking professionals must fully address in order to improve employee productivity, enhance the user experience and solidify their strategic commitment to SaaS and a cloud-first strategy for delivering IT services.

To properly achieve that goal, decision-makers must embrace an enlightened, expanded and modernized view of the role of WAN optimization and application acceleration tools in order to get beyond bandwidth issues of the past and address the bigger culprit of today's application performance: latency. Organizations that rely on SD-WAN alone to improve application performance will find that the latency issue will not be fully resolved.

IT and networking professionals should take a fresh look at WAN optimization tools, even if they have utilized WAN optimization in the past. The dramatic new

demands for improved application performance are only likely to expand in concert with organizations' expansion of both the sheer number of vital applications they must support and the need to have a better application experience in a highly collaborative, real-time decision-making environment.

Riverbed offers IT and networking professionals a full array of solutions in this area, including easy-to-deploy services for enterprise SaaS (including Office 365, Salesforce, ServiceNow and other popular SaaS apps) application performance optimization and mobile client consumption.

For more information about Riverbed's broad and deep solutions set for modern application performance challenges, please visit www.riverbed.com.

¹ "Employees Are Accessing More and More Business Apps, Study Finds," The Wall Street Journal, Feb. 7, 2019

² "Frost & Sullivan Announces Findings of Global SD-WAN Survey," Intelligent CIO, March 13, 2019



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 organizations can deliver a seamless digital experience and drive enterprise performance. Riverbed offers two industry-leading portfolios: Alluvio by Riverbed, a differentiated Unified Observability portfolio that unifies data, insights, and actions across IT, so customers can deliver seamless, secure digital experiences; and Riverbed Acceleration, providing fast, agile, secure acceleration of any app, over any network, to users anywhere. Together with our thousands of partners, and market-leading customers globally – including 95% of the FORTUNE 100 –, we empower every click, every digital experience. Riverbed. Empower the Experience. Learn more at riverbed.com.