

Delivering Application Performance for the Hybrid Enterprise



Hybrid environments present challenges to application performance that haven't existed in the past. Here, discover what those challenges are and how to overcome them.

By Nick Cavallancia

riverbed

Today's users are working from headquarters, in branch offices, remotely, and at home using an increasing number and kinds of devices over combinations of hybrid networks.

You don't really know how all your applications are performing, do you? Given the complexity found in most organizations today, keeping track of which applications are running where, and whether they are functioning at peak performance is a challenge – and you're not alone.

Today's enterprises are a solution-driven mix of corporate and cloud-based data centers with services hosted by public cloud infrastructure providers such as Amazon Web Services, Microsoft Azure, and VMware vCloud Air, as well as SaaS applications such as Salesforce and Office 365. This hybrid enterprise has evolved from a necessity to meet application performance, availability and storage needs for users no matter where they are located. This is why organizations like yours are evaluating each and every application in use to identify the best way to ensure the highest levels of application performance, speed of delivery, and usability for users.

Today's users are working from headquarters, in branch offices, remotely, and at home using an increasing number and kinds of devices over combinations of hybrid networks – MPLS, private VPN, and the increasingly reliable and cost-effective public internet.

And, as if the applications and locations, client types, form factors, and network choices weren't enough, add the fact that users are all expecting an "in-house" or local application experience with performance rivaling that of being in the same building as the application itself – and you've got yourself one very tough challenge.

When it comes to meeting the demand, many types of IT professionals are involved. Those on the application side of the house are concerned about where applications are and performance metrics, while those on the network side of the house are focused on optimizing performance and bandwidth availability while keeping hardware costs as low as possible.

Both groups have the best of intentions, but when critical applications do fail to be delivered and it's obvious something's not working, with so many elements working in tandem, it's difficult to know where the problem is – and who to call to fix it.

Enterprises are no longer supporting just the simple “runs on a single server” applications of yesterday.

So, how do organizations deliver on-premises application performance in a hybrid enterprise with a mix of cloud and corporate datacenters ?

To answer this, let's identify four challenges to application performance that enterprises face, and then look at the steps needed to properly monitor applications in an effort to deliver the performance needed.

Delivery Challenge #1 – Loss of Control

In the days of the corporate data center, it was simple to know where your applications resided, which servers they ran on, how clients accessed them, what the response times were and whether the applications were working. IT had full control over every device and determined which clients could access those devices.

Today's enterprises have lost control of and visibility into applications as they have moved to the cloud.. When applications are delivered by SaaS service providers, IT no longer controls the servers are, or which applications are running on those servers, making ownership of application performance even more difficult. After all, how can you measure (and, hopefully, improve) application performance with reduced bandwidth and latency if you have no control over the application and infrastructure.

And enterprises are no longer supporting just the simple “runs on a single server” applications of yesterday. Today's applications utilize distributed resources, with a set of front-end servers in one location, back-end services in another location and perhaps the data somewhere else completely different. But even with these two or three disparate elements orchestrating an application, the expectation is applications need to run like they are local to the user.

Enterprises also are taking a “whatever the employee wants to use today” approach.. With users taking advantage of their mobile devices, cellular connectivity and anytime, anywhere access to company resources, the complexity of both measuring and delivering application performance is increasing

Users are no longer solely looking to IT to deliver applications; they are taking matters into their own hands.

The problem with all this? You can't improve the performance of an environment you do not control or own.

Delivering Performance: Know What You Have

To be able to deliver on-premises performance, enterprises first need to regain control by establishing not only an application inventory of sorts, but also a true understanding of where their applications reside (if possible, down to the server level) regardless of if the applications are in a local data center, are cloud based, or are SaaS applications. In addition, the connectivity speeds and quality of service guarantees need to be identified for each application to better understand what you have to work with. Without this foundational understanding, enterprises will have little ability to positively impact application performance.

The complexity of today's application environments, mixed with IT's lack of control to deliver a high quality application to users, has created an environment where IT is no longer responding to user needs as quickly as they once have. Because of this, users are no longer solely looking to IT to deliver applications; they are taking matters into their own hands.

Delivery Challenge #2 – Application Sprawl

Enterprises are now consuming more applications than ever before. In addition to applications sanctioned by IT, users themselves are turning to self-service SaaS applications that make their lives easier.

Take the simple example of cloud file sharing – when two users in different locations need to work on a large presentation file and IT doesn't support email attachments larger than 10MB, the users take matters into their own hands and utilize cloud-based file sharing to address the need.

By using these types of applications, the enterprise is now made up of services in unknown locations, self-managed by users, without regard to ease of access or support ability by IT.

This application sprawl puts IT in a position where their users are utilizing more applications than are supported, increasing the complexity of what truly makes up the enterprise. In addition, this movement to utilize any application that helps the user get their job

done only exacerbates IT's lack of control over applications. Lastly, users are consuming unaccounted bandwidth, resulting in a reduction of overall performance of enterprise applications.

You can't improve performance if everyone's doing their own thing.

Delivering Performance: Reel In Rogue Apps

Put simply, application sprawl is a significant threat to IT's ability to reduce bandwidth and hardware costs, as well as their ability to increase performance. With IT getting a handle on where their own applications are, it becomes critical to understand what applications are being used outside of IT and determine whether IT can take control, or provide an approved alternative.

Application sprawl is a significant threat to IT's ability to reduce bandwidth and hardware costs.

Given this sprawl in applications, location, ownership, and complexity, enterprises have less ability to quickly identify the source of a performance problem.

Delivery Challenge #3 - Lack of Visibility

When a user calls the support desk and says they can't access an application or that it's running too slowly, it's no longer an easy answer of "let me see if the server is up." In fact, identifying the problem is the problem! Support may be connecting to the application differently than a remote user, making it a challenge to even see the symptoms users are experiencing. Users may have different devices with varying form factors that a static support desk cannot keep up with. And those form factors may determine which front-end server is utilized, which may differ from the support desk, making things even more difficult.

Now, having control over where your applications are, as well as what applications are in use, are the first steps. It then becomes a question of having visibility into the application, knowing which applications take up what bandwidth, what latency will exist over which paths, etc.

It's this intimate level of understanding that will allow your enterprise to know what is an acceptable performance (from the standpoint of bandwidth, latency and packet loss) on a per-application basis, and what is not. And based on those standards, policies can be crafted

**There are no good
“do it yourself”
answers once your
applications have
gone cloud or SaaS.**

to address issues when user demand of bandwidth from either a business, or non-business application goes outside of acceptable levels and impacts the business.

Take the example of a user streaming video of the Olympics at work. In a perfect world, you'd have policies in place to determine what should be done in this situation so steps are automatically taken to reduce bandwidth utilization.

But when you take this concept out into the cloud along with the remote users accessing applications in the cloud, it becomes even more difficult. Visibility takes on a new meaning when you are relying on uptime and any performance data from the cloud or SaaS provider.

And now you have limited or no visibility.

Delivering Performance: Gain Visibility

It just makes sense: to have true visibility, you need to be able to monitor the application at both the point the user enters the network (or as close to it) and the point they reach the application. That's a tall order to meet, give you have little influence on, say, the folks managing Office 365.

There are no good “do it yourself” answers once your applications have gone cloud or SaaS. So the answer lies in looking to solutions that can be situated near both ends of the path from client to application, being aware of the specific application being used, and intelligently monitoring and measuring the performance of the application.

Once you have this level of visibility and are able to determine when applications are performing, your focus can turn from worrying about bandwidth, latency, and packet delivery, to the more important aspect of application performance that actually drives your business: user productivity.

Delivery Challenge #4 – Productivity

Even given the lack of control over applications and the sprawl in self-service application use, IT is still somehow expected to provide a quality experience for the user. And it's this quality experience that

drives user productivity, and applications performing as if they are local to the user is the real challenge. Just making applications available anywhere on any device alone won't make an employee more productive in the long run.

But we all know it's not as simple as that.

Your entire environment has become dynamic; users can be anywhere, and the server, in essence, can be too (since you no longer have control over where services are being delivered from). Now bandwidth, latency and packet loss become critical factors that determine whether a user is productive or not.

When moving from an in-house or local application to one based in the cloud, bandwidth differences become extremely evident.

The move to the cloud was supposed to fix all your problems – a shift from capital expenditures to operating expenses, fixed monthly charges, uptime guarantees and anytime anywhere access. But without productivity increases, the move is pointless. It's a legitimate concern – when moving from an in-house or local application to one based in the cloud, bandwidth differences become extremely evident. Download times can increase by material factors of 10, 20, or 30+ times that of local performance.

And because client devices – workstations, laptops, mobile devices, and tablets – are now everywhere, connectivity between the client and application runs the gambit with varying levels of speed, reliability, and availability. And yet, they all have the same expectation of a user experience like that of a user sitting in an office located where the application resides.

To be truly productive, your applications must come together across the various networks the data must traverse from front-end to back-end, and present themselves on any needed client form factor to a workforce that is working and mobile.

Delivering Performance: Increase Productivity

In a simple network environment, where you have complete control, it's easy enough to define the quality of service for a particular application. But in the real-world, full of all client, application, and connectivity complexities, there is little you can do without involving a third-party solution.

The move to cloud and SaaS have improved application accessibility, but reduced IT's control and visibility.

The same solution that increased your end-to-end visibility into application performance can be used to focus on how applications impact productivity. Data and metrics turn into useful data that can help you achieve “cloud performance nirvana” much like the on prem user experience...

Conclusion

Words like complex or perhaps even convoluted apply to the environment that makes up your users accessing your applications. It's no longer a straight path from end to end, making application performance a challenging goal to achieve. The move to cloud and SaaS have improved application accessibility, but reduced IT's control and visibility. The shift to enabling remote users with a wide variety of devices has improved usability, but has reduced application responsiveness and lowered productivity.

It's only when enterprises embrace solutions that bring visibility, control and performance optimization to cloud and SaaS applications at levels matching that of local applications, that you will be able to see how your applications are running, how users utilize those applications, determine when exceptions occur, and actually be able to do something about it. ■

With nearly 20 years of enterprise IT experience, Nick Cavallancia is an accomplished consultant, speaker, trainer, writer, and columnist and has achieved certifications including MCSE, MCT, MCNE and MCNI. He has authored, co-authored and contributed to over a dozen books on Windows, Active Directory, Exchange and other Microsoft technologies. He has spoken at conferences such as the Microsoft Exchange Conference, TechEd, Exchange Connections, and on countless webinars and at tradeshow around the world.
