WHITE PAPER

Monitoring the Cloud: Keeping an Eye on What You Can’t See

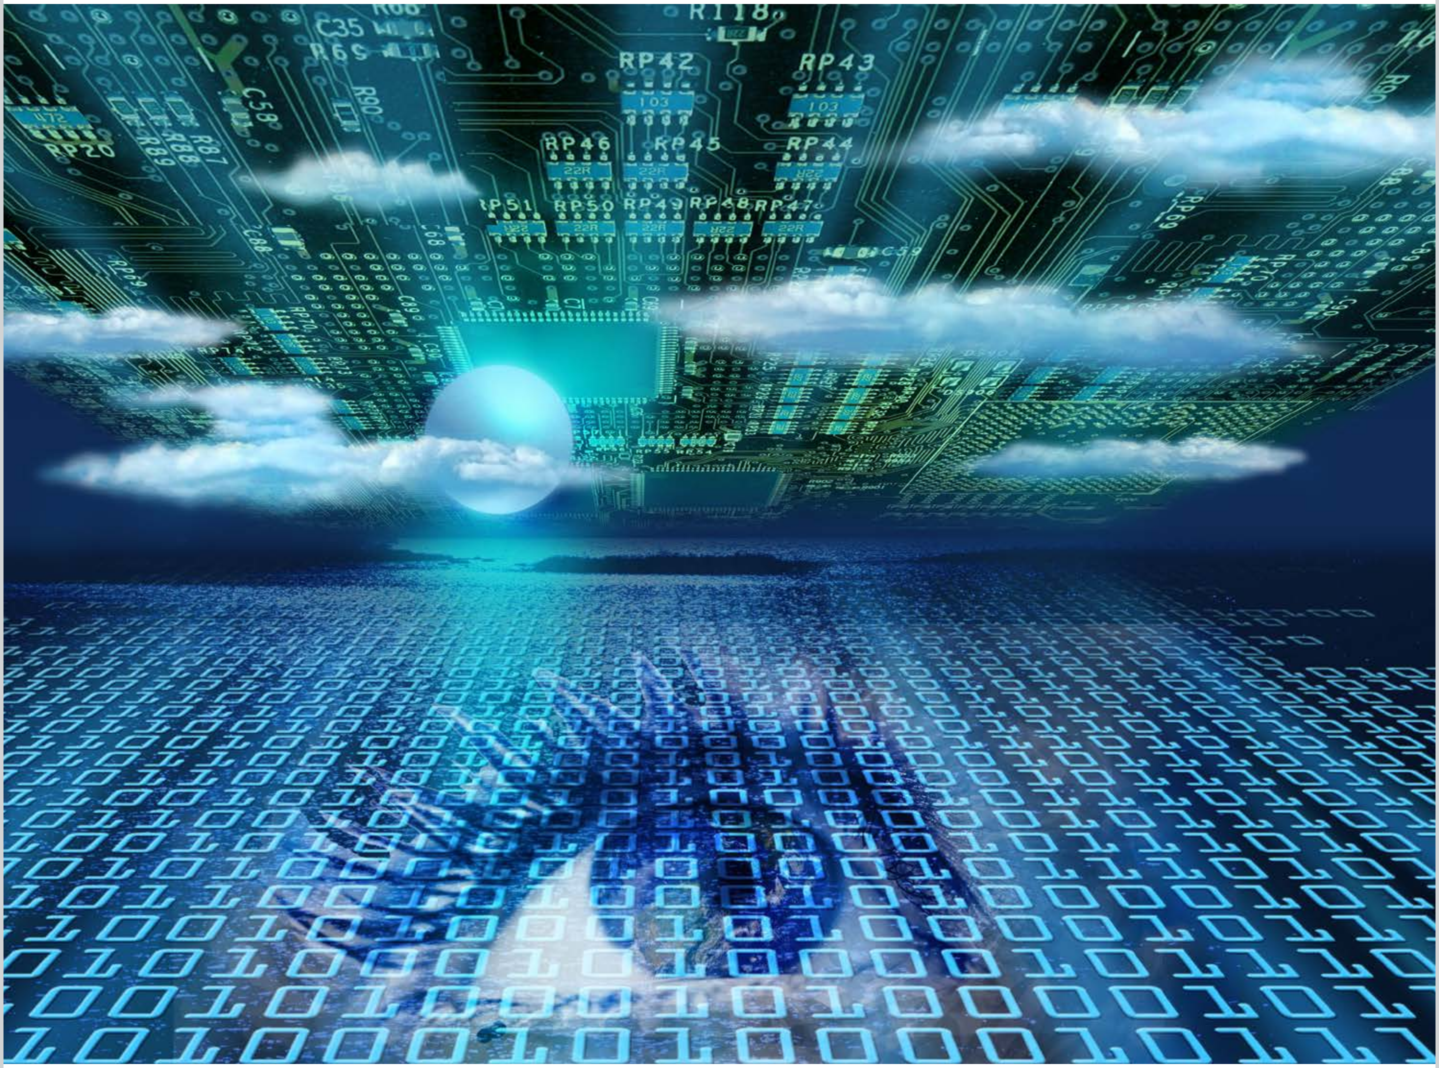

By Nick Cavalancia

PAESSLER
the network monitoring company

TABLE OF CONTENTS

PAESSLER
the network monitoring company

Before the cloud, monitoring was pretty clear. You watched an environment that you owned and managed. It was your servers, your network, your routers, your switches, all of which gave you broader visibility into whether an issue was isolated to a single server or to the entire network.

And with the move to the cloud, some of that visibility gets hazy. Slowdowns in application performance can't easily be attributed to your network, the connection to the cloud, or the servers hosting the application. Analyzing the cause of an issue requires calls to an ISP support team that may or may not have the same sense of urgency as you, because they're, technically, meeting their SLA.

So, is monitoring even valuable, when it comes to cloud applications?

Monitoring is only as good as the action you take because of it.

Monitoring is only as good as the action you take because of it. If you only know your application is running slower than normal, but have no actionable intelligence of where the problem may reside other than a red light on a dashboard, the monitoring wasn't all that helpful and, somewhat, useless. What you really need—ever more so once you've moved to the cloud—is lots of more specific detail that's truly actionable intelligence around the issue at hand. That way, you can take action—even when it's an issue with your cloud provider—making monitoring valuable again.

But, by definition, moving an application from on-premises to the cloud changes things. There are new unknowns around the server, connectivity, response time, and more, because they're all "out there" and no longer under IT's control. And because you want to ensure users aren't even aware they're going to the cloud (that is, you strive to maintain on-prem-like performance), monitoring becomes all that much more important.

So, how do you monitor the cloud in a way that both gives you the actionable information that you need as well as takes into account the unknowns of the cloud?

In this whitepaper, we'll look at some of the changes moving to the cloud has on monitoring, and look at what performance factors you need to focus on to ensure you maintain not just visibility into the performance of your cloud-based applications, but also obtain actionable intelligence to quickly respond to issues.

Monitoring Forecast? A Bit Cloudy

It's pretty simple when it comes to on-premises monitoring. You include devices on the network, your applications, servers, disk, CPU, and network throughput, etc., place them up on a dashboard, setup some alerts, and you have an ability to see when some part of the networks not performing well, and can drill down into the details to identify why.

But when you add the cloud into the conversation, it's no longer that simple. There are aspects of the environment that change that impact your ability to monitor applications in the same manner as on-premises, and leave you lacking in a few areas:

There are aspects of the environment that change that impact your ability to monitor applications.

Lack of Visibility

When something is going wrong with an application in the cloud, the biggest challenge is not knowing there's a problem, but where in between you and the cloud is the problem. Let's say your Exchange Online within Office 365 is down. You have visibility into the fact that it's down, but you're not exactly certain where the problem is, because the cloud introduces a large number of unknowns. It is just you locally? Is it something with your connection to the Internet? Is it something wrong with the Internet itself, just routing on the Internet? DNS? Perhaps Office 365 is experiencing a denial of service attack. There are so many possibilities, so troubleshooting the problem is going to require a lot more investigation and a lot more of your time to figure that out

Lack of Ownership

The use of the cloud for applications really falls into two camps—software as a service (SaaS) or infrastructure as a service (IaaS). In both cases, you no longer own any of the servers, networking hardware, or firewalls. And in the case of SaaS, you also no longer own the operating systems or applications. You've exchanged all that for an easy monthly payment and an uptime SLA. So gaining access to monitor these parts of your environment becomes a challenge.

Lack of Control

It used to be once you identified a problem (and have enough visibility to determine what the specific issue was) you'd simply pick up a phone and call the network person who's just down the hall from you to ask them to address the issue immediately. With the cloud, you're nothing more than a number in a queue with an SLA.

And it's these deficiencies impact monitoring, making it (for those of you still thinking of moving to the cloud) more difficult for you to move there, or frustrating for those of you already there. Most of what's been discussed so far is all conceptual, and the move to the cloud has some very real implications on what, how, and whether you can monitor.

So, how does the cloud specifically complicate monitoring?

Moving to the cloud has added some complexity to your network, with the biggest change being the path from here to there.

The Cloud's Impact on a 360° View

Moving to the cloud has added some complexity to your network, with the biggest change being the path from here to there. It's somewhat like having a really long Ethernet cable (without the attenuation issues). There's also the dependency on the provider, whether servers are virtualized, and does the cloud environment have enough resources. It can easily get pretty confusing in terms of where to place your focus in trying to obtain a 360 view of this new extended network.

So, what are the ways the cloud impacts monitoring?

Latency

When on-premises, latency wasn't an issue. The distance between your users and the applications they need was maybe one or two hops on a switch within the same building. But once you've moved to the cloud, latency become much more an issue as you now add additional hops to the route just to get to the cloud provider, more distance to travel within the cloud provider's network, whatever response time the application in question needs, and then traversing the same path in the other direction.

Moving to the cloud simply takes the user experience and adds more delay onto it. Now, in the network sense, we're talking about delays in terms of milliseconds, but latency ends up being a multiplier for any issues that arise on the backend as well. For example, a hosted application seeing a decrease in performance over a latent connection will appear even slower to the end user. And the further you stretch that "Ethernet cable" the more hops, the more latency, and the more multiplying that occurs.

And when you're trying to monitor, latency makes it a bit difficult to know whether something actually a problem, or if an application is just slow "because of the cloud".

Moving to the cloud simply takes the user experience and adds more delay onto it.

Application Performance

When you move applications to the cloud, you expect that they should run as well or better than they did on-premises. But, with the potential lack of control over what resources are devoted to a given application, as well as the change in network architecture at the cloud provider, you can't simply assume performance will remain identical to what it was on-premises.

There are a number of cloud factors that impact application performance. Hardware is the obvious one. You run your applications on-premises using a specific level of hardware, but when in the cloud, something a bit more robust is required in order to cut down on the potential processing and disk latency that can exist, that would then be multiplied by network latency. And if you're moving from physical on-prem to virtual in the cloud, it's important to recognize that not every application—especially those that are IO intensive—runs as well when virtualized.

Beyond just hardware, other cloud factors impact your view of the problem: the cloud provider's local network, the impact of latency on client requests and responses, and the latency between tiers of a multi-tiered application are just a few.

Environmental Issues

Even when in the cloud, we're still talking about servers in a data center or two somewhere. And even that data center can experience issues around weather. When the application was on-premises, you know whether your having, say, lightning storms or a hurricane (you can just look out the window). But, in moving to the cloud, environment conditions plat the very same role on a data center's availability—just you aren't able to know about those conditions as easily.

Even with the move to the cloud impacting a comprehensive view of what's occurring, there are areas you can focus your monitoring energies on to ensure you have the highest availability and performance possible from the cloud.

There are areas you can focus your monitoring energies on to ensure you have the highest availability and performance possible from the cloud.

Cloud Performance: Where to Watch

When you think about the performance factors that you need to monitor because you're watching the cloud, the first thing that comes to mind is that, in some ways, nothing changes. You still have to worry about many of the same factors—application performance, services running, server hardware performance, virtual host performance, disk, CPU, memory, NIC, and bandwidth.

But, at the same time, everything changes, especially from a perspective of base lining performance. With the latency multiplier effect, that application response time of two seconds you achieved on-premises may be higher when that same application is in the cloud. That means, as you move an application to the cloud, you're going to have to level set just about every measurement in order to define what is and isn't acceptable performance.

So, what should you be monitoring?

The Path from Data Center to Cloud

The path taken will depend on the application used. For some, the only think on-prem is clients. For others, there may be an application with a dedicated server solely responsible for communicating with the cloud. And still for others, the path may a number of applications, mixing and matching the previous two scenarios. So, you need to identify exactly what's in the cloud that you're talking to, and who on-prem is communicating with the cloud.

You should be watching each part of the path, ensuring devices are functioning and running within acceptable performance specifications, including:

- Gateway devices
- WAN interface
- LAN traffic
- Your ISP

Additionally, watching values around jitter, packet delay variation, roundtrip time, as well as the existence of lost, duplicated, corrupted, or out of order packets for each step of the way will also provide you context around where problems may lie.

It's obvious you're taking some risk of increasing your latency by moving to the cloud.

Latency

From the previous definition, it's obvious you're taking some risk of increasing your latency by moving to the cloud. To address this, monitoring the cloud provider at a few points in the path, measuring jitter and response time, will allow you to both identify latency issues, as well as pinpoint where problems originate. These points include:

- **End-to-end**—monitor from your LAN to the application to get an overall view of latency.
- **From your application server to the cloud provider's gateway**—identifies whether or not the cloud provider's LAN environment is adding to the latency issue.
- **QoS from your data center to the cloud provider's gateway**—allows you to identify if the issue lies in transit across the Internet.

The Application Environment

When you move an application to the cloud, you're paying for service from a cloud provider. And that service comes with some expectations of performance and availability—expectations you can monitor once implemented to ensure your expectations are being met.

Depending on what kind of environment your applications reside in, you are looking to have a robust backplane, redundancy to the nth degree, and ample hardware resources for the server or VM. Once those things are established, you monitor those specifics that can be (such as hardware resource usage, network utilization, etc.), and for those that can't (such as the backplane) you monitor the overall user experience, beginning with a focus on response time.

And don't forget about multi-tiered applications; you need to including monitoring of both each server involved, as well as (to the best of your ability) the specific communications and channels used between those servers.

Virtualization is a staple in the cloud playbook.

The Virtual Environment

Virtualization is a staple in the cloud playbook. You've probably moved to a virtual environment for reasons like a smaller hardware footprint, lower power needs, ease of backup/replication or management of the entire system, etc. While virtualization does simplify a lot, in the world of monitoring, the move to virtualization yields additional layers to potentially create problems.

Your server now is just one of many on the same hardware, and when issues arise with an application or an entire VM, you now need to also be looking to see if any part of the virtual infrastructure—for example storage and RAM provisioning—is playing a role.

On-Premises

You, of course, still have to monitor your own physical environment in the same way you always have, maintaining two sets of performance baselines. All of the "is it the local network" parts of isolating the source of a cloud application issue would fall under here.

Seeing Through the Clouds

Even with moving applications to the cloud, the goal for monitoring is still the same—maintain detailed visibility into your critical infrastructure, systems, and applications, along with the interactions between each them and their clients.

So, really, the move to the cloud doesn't require a shift in your monitoring strategy; it just needs some changes in execution. The breadth of devices, networks, and layers of sensors necessary will increase, as will the detail required. But in the end, providing you have some means of monitoring access to your cloud providers network, hypervisors, and applications, you can see through the clouds and achieve on-premises-like visibility. ■

The move to the cloud doesn't require a shift in your monitoring strategy; it just needs some changes in execution.

With nearly 20 years of enterprise IT experience, Nick Cavalancia 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 tradeshows around the world.