



Visual Studio  
Magazine



# The Cloud Application Performance Management Paradox

Application Performance Management might not be top of mind for companies moving to the cloud, but it should be. Here's how to handle APM in the cloud.

BY RICHARD SEELEY




# One source of truth


See all your data. Boost performance. Drive accountability for everyone.




**Mobile Developers**  
End-to-end visibility,  
24/7 alerting, and  
crash analysis.



**Front-end Developers**  
Deep insights into  
your browser-side  
app's engine.



**IT Operations**  
Faster delivery.  
Fewer bottlenecks.  
More stability.



**App Owners**  
Track engagement.  
Pinpoint issues.  
Optimize usability.

Move from finger-pointing blame to data-driven accountability.  
Find the truth with a single source of data from multiple views.

[newrelic.com/truth](http://newrelic.com/truth)

**Major cloud vendors, including Amazon, Microsoft and Google, make claims for 24/7 availability at high performance levels, and the data backs them up.**

## Introduction

Application Performance Management (APM) for the cloud is not fundamentally different from what APM was in the days of client/server, glass rooms and data centers filled with mainframes, say analysts who cover the issue. As a recent [article](#) in Fortune pointed out, “IBM often expresses the notion that the mainframe is the ‘original’ cloud,” analyst Charles King of Pund-IT is quoted as saying. He notes that “... the management tools that are key to cloud computing originated decades ago in mainframe systems.” Beyond the hype, the cloud is just another computing platform and the basics of what end users expect in terms of application performance have not changed regardless of whether they were being accessed from the “dumb terminals” of the 1970s or iPhones and iPads today.

However, there is a paradox where APM for the cloud is the same as APM for applications running on-premises, and yet it is also new and unique. The goal is for the business user, partner or customer to have a quality experience working with the app whether they are accessing it from a PC, laptop, tablet or smartphone. Most of the old APM rules apply. Clunky code is still often the culprit. The platform, not so much.

Major cloud vendors, including Amazon, Microsoft and Google, make claims for 24/7 availability at high performance levels, and the data backs them up. Analysts agree that cloud providers are doing a good, often near-perfect job.

“The cloud is actually turning in some good uptime figures,” says David Linthicum, consultant and author of [Cloud Computing and SOA Convergence in Your Enterprise: A Step-by-Step Guide](#). “If you look at [Cloud Harmony](#), which is the site that tracks cloud outages, the track record has been pretty good, relative to even internal systems. So, they seem to be providing better uptime today and better reliability today than most internal systems are able to do.”

Cloud providers offer service level agreement (SLA) compliance that runs as high as 99.999 percent, which is sometimes referred to as “100 percent SLA with a 10-minute exclusion,” according to Gartner’s [Magic Quadrant for Cloud Infrastructure as a Service](#). Monthly availability for SLAs at 99.95 percent or higher tends to be the norm. This is typically better than SLAs for managed hosting, according to Gartner. High cloud compute availability is a product of “redundant infrastructure in conjunction with virtual machine (VM) clustering, or the ability to rapidly detect VM failure and immediately restart it on different hardware,” Gartner reports.

**The major difference in cloud versus traditional computer room platforms is the number of moving parts in the cloud.**

Data showing this SLA track record for the cloud infrastructure makes it difficult to blame the platform when a customer abandons a shopping cart or a business user phones to say the accounting app is running way too sloooooow.

Poor performance of apps results in lost business. In a world where customers are accessing apps from social media this is critical. Statistics from KISSMetrics indicates that 25 percent of users abandon a web page if it takes more than four seconds to load. And businesses may not get a second chance. That same study found that after experiencing poor performance, 86 percent of users delete the offending application.

The major difference in cloud versus traditional computer room platforms is the number of moving parts in the cloud. Researchers at Gartner call this the “unique challenge for application performance monitoring.” This is because the applications are often Web services running on various virtual systems most of the key system components are no longer hosted on a single machine. “The applications themselves are very likely to be moving from one system to another to meet service-level objectives and deal with momentary outages,” according to Gartner.

In that environment, how is IT going to monitor and manage cloud applications?

Analysts who research cloud technology say the APM tools available from the cloud vendors and from third parties are pretty good. Of course, those tools can't help if IT departments don't use them. And that is often a primary problem. It's the old story: first you need to plug it in. There is no magic help for organizations that ignore APM and expect apps deployed to the cloud to run smoothly on their own in some unknown environment. One possible problem with APM for the cloud may be that too few people are paying attention to it. This white paper seeks to remedy that by exploring what APM for the cloud can do for you and what you need to do to get the most out of APM and the cloud.

### **What APM for cloud means**

How do analysts define application performance management in cloud environments?

A straightforward answer comes from author and consultant David Linthicum: APM focuses on everything having to do with the ability to deliver application services out of public and private cloud environments. It covers

**A recent survey of organizations venturing into the cloud indicated that APM standards need to be the same regardless of where the application is running.**

the end users and all the points in between, network application, database, compute, scalability, and the ability to work efficiently during saturation times.

A recent survey of organizations venturing into the cloud indicated that APM standards need to be the same regardless of where the application is running, said Joe McKendrick, an independent analyst, who conducted the survey. This is especially germane as many organizations will be taking the hybrid cloud approach with legacy transactional applications running in an on-prem computer room while apps for mobile and social media customers are coming out of the public cloud.

“Application performance management should be the same whether delivered from an on-premises data center, or from the cloud,” McKendrick said. “APM is concerned with delivering the most amount of uptime, and the fastest and most responsive performance, to end users—regardless of whether they are accessing applications from the cloud, or from within the enterprise.”

In the hybrid cloud world, APM needs to be working everywhere, opines Jason Bloomberg, president, Intellyx LLC, who covers cloud APM vendors for clients. “Perhaps there’s a CDN (Content Delivery Network) involved, but we can consider those to be part of the cloud. True, the on-prem back-end systems may not be in a cloud, but any cloud APM product worth its salt will need to provide some level of visibility into the performance of such systems, for those applications that leverage systems of record.”

### **Understanding the cloud APM paradox**

How is APM in cloud environments different from traditional data centers? The short answer is that the public cloud is out of your control. Prior to the advent of the cloud, the IT department was in control of the entire environment.

“With traditional data centers, you’re in charge of everything,” Linthicum explains: “You own the platforms, you own the networks, you own the transfer, the data, the databases, the applications, and from the applications to the end users and all points in between.”

But if you migrate one of your applications to the cloud, it is now running somewhere on platforms controlled by the cloud vendor, whether it be Amazon, Google, Microsoft or whoever. Your IT department is no longer in charge of managing the platform, managing the environment, managing scalability, managing load distribution, managing failover, managing the network, even managing the database.

**70 percent of application managers in “distressed environments” on premises are anxious to take advantage of cloud scalability.**

Some IT managers may find this prospect a little scary. But analysts say it doesn't have to be if IT uses cloud APM tools properly. Keeping their eyes on the APM dashboard, an organization's in-house IT staff can make sure their apps in the cloud are performing at expected levels with consistent high availability. Application managers will also be able to make sure they are achieving the key benefits of the cloud, including minimizing if not eliminating downtime, and maximizing performance while benefiting from scalable processing power that is likely to far exceed what would be available running on-prem.

Cloud application performance is seen as an advantage by many of the 741 application managers in the survey McKendrick recently conducted for Unisphere Research ([“Performance Under Pressure: 2014 OAUG Survey on Enterprise Application Management”](#)).

Far from perceiving cloud application performance as a problem, 70 percent of application managers in what the analyst characterized as “distressed environments” on premises, are anxious to take advantage of cloud scalability.

“The survey shows that private, hybrid and public clouds are seen as a way by many enterprises to achieve their availability and performance goals,” McKendrick explained. “Cloud-based solutions—whether from public services, private cloud, or a hybrid of public and private approaches—are seen as offering greater capacity on demand, as well as behind-the-scenes tuning expertise, to help address performance issues.”

The application performance advantages of the cloud are resulting in more IT departments migrating or planning to migrate apps to private, public and hybrid clouds, according to the survey. At least 35 percent of those surveyed plan to migrate some applications to cloud environments provided by third-party vendors. The analyst predicts that percentage will double in the next three years. Within five years, he expects almost all IT organizations will have applications in a private cloud, at the very least. Those on the leading edge are already in the cloud.

Given the anticipated level of cloud migration and the performance expectations, there will be a growing need for APM tools to make sure that goals are being met.

**The cloud may become so ubiquitous that the term itself will disappear as it reaches the hype cycle peak.**

### **Coping with the coming cloud disruption**

While most IT managers plan to migrate at least some applications to the cloud, few of them are sure what they are getting into and the only sure thing is the disruptive nature of the technology. The cloud may become so ubiquitous that the term itself will disappear as it reaches the hype cycle peak. The bleeding edge technology of 2008, may just become the basic compute environment, some analysts predict.

But between now and then there may be a seismic disruption where computer terminology may go topsy-turvy.

“Today’s digital environments upend existing market categorizations,” says Bloomberg, “so the notion of “cloud APM” is becoming increasingly incoherent, as both ‘cloud’ and ‘application’ are terms with shifting meanings.”

“Fundamentally, today’s applications are end-to-end, comprising web, mobile, and perhaps other touch points on the frontend, a combination of SaaS, PaaS, and IaaS, and on-prem systems behind the scenes,” the analyst says. “All elements of this complex, multi-part interaction must deliver high performance and scalability, and it falls to today’s APM tools to ensure the performance meets the business goals for the application.”

As an example of the complexity, Bloomberg says you only have to look at a simple Web application. One Web page might contain dozens of widgets or other third-party services, including analytics widgets, and dynamic advertising. The page itself might be the product of one SaaS app while the various services are running on many other cloud environments. Mobile apps incorporating third-party services are also interacting with the page where users have the expectation of high performance. With a mobile airline app there is interaction with systems of record that may be running on an old reliable mainframe. But frequent flyers using smart phones expect near instant access to their reservation data.

Making sure performance expectations are met would appear to be a daunting task. But the important thing to keep in mind is that however complex the cloud may be with all its moving parts, it is still basically a computer platform.

### **The tale of the latency tail**

Determining application performance in the cloud requires new ways of looking at the data on an APM dashboard. For example, judging an

**Looking at good averages on an APM dashboard may not give you a full picture of frustrations some end users are having with an application.**

application's latency issues requires deeper insights. Looking at good averages on an APM dashboard may not give you a full picture of frustrations some end users are having with an application.

This issue was highlighted recently by Google's Brad Abrams in a [GCP Live Session: Debugging the Cloud: Troubleshooting Evolved to the Cloud Era](#).

He points out the 100 milliseconds average response time in his demonstration application is considered "pretty good."

"But averages, when you're talking about the latency of an application can be very misleading," Abrams explained. That is because latency distribution is expected to fall on a standard bell-shaped distribution curve. "But if you've actually done some latency work, you know it rarely looks like this."

Latency often falls on a long declining curve and users unlucky enough to be at the end of that curve are not getting response times anywhere close to 100 milliseconds. They may be getting frustrated with the application's slow response times even though users at the head of the curve are happy. This is what the Google brainiacs call "long-tail latency."

"It's important to examine the long latency tails of systems, even when they appear fast." according to Jeff Dean, a senior fellow at Google, who is quoted in Abrams' presentation. In judging application performance, Dean recommends using percentiles, rather than averages.

While the average response time was 102 milliseconds in the cloud app that Abrams was demonstrating, that does not cover all the users. Even if 90 percent of users were experiencing the faster latency times, looking at the 90th percentile users, they were getting 300 milliseconds response times. Those slower response times may not be acceptable for a mobile app as users unlucky enough to be in the 90th percentile or above may abandon it and never use it again.

Abrams recommends looking at the long latency tails and fine tuning the application to achieve optimum response time for all users.

### **Cloud vendors offer APM tools**

APM Tools are available from the major cloud providers as part of the services offered to customers. [Amazon CloudWatch](#) is a monitoring service for AWS cloud resources and the applications running on AWS. IT can use

**APM Tools are available from the major cloud providers as part of the services offered to customers.**

Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms.

Google has just released Cloud Trace, a feature of the Google Cloud Platform that allows IT to view the RPCs (remote procedure calls) invoked by an App Engine application and to view and analyze the time taken to complete each RPC. It traces all RPCs made by an application and those traces can help identify performance bottlenecks in cloud apps, according to Google.

As part of its Azure cloud offering, Microsoft provides Azure Diagnostics, which captures performance data including application-level log files, as well as custom logging and tracing information. Azure SQL Database Dynamic Management Views is designed to provide IT with feedback for diagnosing performance problems in the Azure SQL Database.

### **APM Tool Vendors**

When consulting with organizations that are planning to migrate applications to the cloud, Linthicum suggests leveraging best-of-breed third party APM tools. He recommends looking at third-party tool providers that have partnered with the cloud vendor. Those tools will be tuned to provide key metrics for applications in that specific environment.

APM tool vendors include:

**AppDynamics** traces transactions end-to-end, provides database monitoring, and real time information on end-user experiences.

**CA** offers an Enterprise Mobility Management suite that focuses on the growing mobile market providing “context-aware management across mobility use cases.”

**Compuware** offers a suite of APM products as part of its Dynatrace brand. It provides “user-centric” APM designed to show IT how applications are performing from the end user’s point of view.

**Dell Cloud Manager** covers deployment and management of enterprise applications across private, public and hybrid clouds. It includes a suite of tools designed for private and public cloud platforms. It is available as SaaS or as on-prem software.

**IBM Application Performance Management (APM)** is designed to identify bottlenecks and determine the root cause of performance problems. IBM says its integrated analytics improve application availability by 60-90

percent and reduce the lengths of outages and slowdowns by 50 percent or more. IBM Application Performance Management tools are available in SaaS and on-prem versions.

**Logic Monitor** touts hassle-free monitoring for APM with alerts sent to key personnel via emails, SMS and phone.

**New Relic** monitoring products include Deployment Reports showing before and after pictures of app's performance following an update so IT can back out of a problematic change before it impacts end users. Transaction Tracing provides code level diagnostics and full stack traces. Cross App Tracing covers transactions across tiers. Integrated alerts appear on one unified dashboard.

**Consultant  
David Linthicum  
recommends looking  
at third-party tool  
providers that have  
partnered with the  
cloud vendor.**

**Right Scale** Cloud Portfolio Management covers applications that may be moved between clouds as business needs evolve. This tool focuses on providing IT with control over their organization's cloud usage to help manage costs.

### **Case Study: APM at Kellogg**

The Kellogg Company, founded in 1898 and famous for its Frosted Flakes, is on the leading edge deploying mission critical applications to the cloud using software from SAP running on Amazon Web Services with APM from Amazon CloudWatch.

In the fast-paced world of digital marketing and TV advertising, Kellogg needs to run dozens of complex data simulations on promotions. By 2013, the traditional applications running on an on-prem platform only had the capacity to do just one simulation a day. That was too slow for today's marketplace and it had a major impact on the company's bottom line. "Any improvements we make go straight to our bottom line," Stover McIlwain, senior director of IT infrastructure engineering at Kellogg in an [Amazon AWS case study](#). "If we improve trade spend by just 1 percent, that's \$50 million dollars."

"Margins are very tight in our industry, and even slight changes in trade spend can swing market share," McIlwain explained. "Revenue growth is flat in some of our categories, so we need to be very agile to stay competitive."

To achieve the agility Kellogg needed, the company selected SAP's Accelerated Trade Promotion Planning software utilizing SAP HANA in-memory database technology platform running on the AWS cloud environment.

For APM, Kellogg uses Amazon CloudWatch to keep the applications running efficiently while processing 16 TB of sales data weekly from promotions in the US, and also modeling dozens of data simulations a day. It is also credited with lowering costs by helping to track and allocate usage.

“CloudWatch helps our people make better decisions around the capacity they need, so that they can avoid waste,” McIlwain told Amazon. “We were never able to do that with our on-premises infrastructure.”

**The Kellogg Company, founded in 1898 and famous for its Frosted Flakes, is on the leading edge deploying mission critical applications to the cloud.**

### **Summary: For the best performance, take it slowly**

Because organizations may encounter application performance issues in cloud adoption, a cloud migration strategy starting with the least time-sensitive applications is what Linthicum recommends to organizations. At first, companies are more likely to be focused on issues like security as they contemplate moving to the cloud. APM is not at top of mind yet. To cope with performance issues a gradual migration strategy may work best. For example, migrate applications used by your own human resources or accounting departments first. If there are performance issues, the end users can wait while IT troubleshoots the problems. Lessons learned in those migrations will help smooth the way for mission critical e-commerce apps where poor performance might result in loss of business and customers. It also may hasten the adoption of APM tools.

**ADT**

---

*Rich Seeley is a technology editor specializing in business software development strategies for Cloud, Mobile applications, Service-oriented architecture (SOA), and Web services. He has worked as an editor for Application Development Trends and Campus Technology.*

---

