

Borderless applications and the changing face of application performance management

Application outages and performance glitches plague our daily lives. A quick scan of recent headlines shows why this is concerning:

“Salesforce.com customers buzz on Twitter about nationwide outage, say service restored in about 30 minutes”
“iTunes Store/App Store Outage?”

“Twitter for BlackBerry Experiencing Performance Issues”
“Facebook Sees Major Outage — Takes Out Partner Site Plugins, Too”

“Netsuite Back Online After Cloud Apps Outage”

“Toktumi Line 2: Popular iPhone App Experience Outage, New Sign-ups Shut Down Temporarily”

“The Blackberry outage that RIM does not want to talk about continues to affect people on both sides of the Atlantic with outages reported in the United States”

As applications become more collaborative and are accessed globally, businesses will need to ensure they are well-equipped to manage user performance and eliminate outages in order to compete. Understanding application performance and user experience for these types of transactions is critical. Today, applications have a direct impact on company revenues, customer satisfaction and brand perception. However, being able to manage the business application's performance has become a difficult task for IT operations, due to:

- increasingly complex and distributed application architectures
- application usage that spans globally
- an explosion in the types of devices and browsers.

In order for businesses to thrive and increase market share, they need to be innovative, use cutting-edge technology and market themselves in sophisticated ways. How can critical business applications be managed without increasing budgets and hiring more people? IT operations play an

enormous role in ensuring business success. IT is the “mother ship” for managing the day-to-day business operations. It plays a pivotal role in helping the business change to meet the demands of today's users.

IT managers and their business counterparts need to differentiate themselves from the competition without adding resources, and even producing more with less. Three questions come to mind when seeking solutions to achieve business transformation, competitive differentiation and growth in revenue:

- How does IT align itself to meet business needs?
- How can IT capitalize on what's already in place and innovate without adding cost or overhead?
- How will IT meet customer application performance expectations, maximizing revenue opportunities while safeguarding the corporate brand?



The impact on the business can be great if IT is ill-equipped to ensure performance meets user demands. If we look back in time to see how IT and business owners have operated, we can see how drastically people, process and technology have changed and where they are headed in the future.

As the business delivers new revenue-generating applications to its customers, IT is looking to optimize cost and increase efficiencies through the adoption of new process, new technologies and new tools. Furthermore, ensuring high performance for newly adopted technologies becomes a “must,” due to the threat of competition and loss of revenue. End-to-end application performance management becomes critical to both IT and business counterparts, with the need to track, measure, identify and fix issues prior to users feeling the effects of poor performance. End-to-end visibility must be achieved no matter where the application resides: external, internal or cloud-based.

Therefore, the critical driver for ensuring that performance meets user demands is the ability to manage the application across the entire application delivery chain anywhere around the globe. An end-to-end application performance management solution that can easily extend to managing new technologies provides a foundation for IT and business counterparts to collaborate and address performance issues quickly and meet business demands. This paper discusses the shift in application performance management by IT operations in three key areas:

1. Application architectures
2. Zone of responsibility
3. Customer expectations

In addition, it provides insight into how best to ensure applications spanning across all boundaries of enterprise and the Internet can successfully be monitored and managed to meet the needs of the business. A use-case scenario, based on a fictional banking business, Paramount Funds, will help describe the changing face of Application Performance Management from the past, to the present and future. The financial institution referenced in this paper is fictional, but the use case describes situations common within IT and business organizations in many types of companies around the globe.

	PAST	PRESENT	FUTURE
Overview	Internal-facing application— Business focused on internal productivity	Internal- and external-facing applications— Business focused on revenue growth and brand recognition	External-facing applications— Business continues to focus on revenue growth and brand recognition
Application Architecture	<ul style="list-style-type: none"> - Single data center - 2-3 Tiers + MF - Physical, static, thick client + 1 browser - LAN/WAN 	<ul style="list-style-type: none"> - Multiple data center - Hosted applications, 5-25 Tiers + MF - Physical + virtual, static - multi-browser— Internet 	<ul style="list-style-type: none"> - Cloud computing - Highly virtualized, dynamic, highly mobile, multi-browser— Internet+
Zone of Responsibility	<ul style="list-style-type: none"> - Availability only - Internal data center only - Infrastructure-centric + application - Silos in data center 	<ul style="list-style-type: none"> - Availability expected - Performance, browsers , data center + Internet, CDN, service providers - Application-centric + service - Horizontal across data center + Internet 	<ul style="list-style-type: none"> - High availability required - Mobile performance, + mobile networks CDN, service providers - Optimized service - Cloud
Customer Expectation	<ul style="list-style-type: none"> - Managed by human interaction - Availability OK 	<ul style="list-style-type: none"> - Self-service, simple tasks - Availability expected - Performance matters 	<ul style="list-style-type: none"> - Self-service, all tasks - Across the globe, any time - Performance demanded, instant expectations

THE FACE OF THE PAST

A decade ago, critical applications had simple architectures that were typically served out of a single data center and were under the complete control of the IT organization. The business focused solely on generating revenue through great face-to-face customer service. Critical applications that were used to execute customer transactions had to be available to ensure employees could successfully satisfy customer needs.

PAST APPLICATION ARCHITECTURE

The application's infrastructure was static and the physical components were critical for ensuring the application was available for users to access information. Those physical components consisted of a mainframe, database and application server—a typical two- or three-tiered environment. The monitoring or management of these physical components was on a siloed team basis, where one monitoring tool managed that team's portion of the application infrastructure. The goals of the IT operations organization were simple—keep the data available—while application performance was a secondary concern. More often, the way to solve performance issues was to buy more servers and place them on the network. Poof, performance problem solved; however, IT budgets became bloated.

If the application became unavailable, IT operations would call each team into the “war room” to strategize on where the problem resided. To solve the application problem took much time and effort and cost a lot.



Internet applications from the past, such as this one from Compuware Corporation, dating back to 1996, are extremely simple, using very few visuals to help draw attention or entice users to click around at different options.

PAST ZONE OF RESPONSIBILITY

IT operations had individual groups that managed components of the applications. Each IT group had its own monitoring tool to understand whether its portion of the application infrastructure was available and performing well. No group had a need or desire to communicate how performance might have impacts on the other IT group components, let alone on end-user productivity. This method of managing the business's mission-critical applications was successful due to the lower expectations of the end users. End users (employees or business users) didn't know or expect the application to perform any better than it did. The only expectation the users had was that the application would be available. In some cases, users could still get the job done by manual means, but over a longer period of time. But when the application continued to be unavailable, the effects on the business became greater in cost, loss of productivity and lost revenue.

In many aspects, IT operations focused solely on whether each group—database administration team, application team, network team, etc.—had data available, which was fine for the business at the time. Business reliance on the application was starting to increase, but as long as IT could only report on application availability, there were no other indicators that would help the business take other steps for improvement.

PAST CUSTOMER EXPECTATIONS

Customer expectations in the past were mainly focused on how well they were serviced. The business users who serviced customers focused solely on how well they could interact and provide meaningful service that kept customers satisfied. Most of the customer interactions with business users were direct, with customers going to a physical business location to perform a transaction. The face-to-face customer service interaction was the means of doing business. The more personable and comfortable the interaction was between business users and customers, the more likely the overall goal of making a profit and growing revenue.

SCENARIO 1: HOW PARAMOUNT FUNDS TACKLED APPLICATION ISSUES

The Paramount Funds customer was used to driving to the location, waiting in line to perform a transaction and leaving the branch happy. On this particular day, the Paramount Funds location in Boston was experiencing application availability issues. The tellers couldn't get access to the “withdrawal” function for customer accounts. A temporary work-around was put into action, as the tellers started to perform withdrawals manually until the withdrawal system became available. In the meantime, customers were negatively impacted when the lines inside the bank and outside the drive-up tellers were growing longer by the minute.

The Paramount Funds manager called the IT help desk to tell them, “We have a problem and it’s critical. Our customer account withdrawal system isn’t available to perform any transactions for our customers.”

At this time, the help desk submitted a priority one call ticket for all IT operation teams to solve the problem, since the help desk team was unaware of where the problem resided. The network, server, database and application teams were all called into a room where IT management said, “We have a problem. The Boston Branch can’t perform ‘withdrawal transactions’ for their customers and we need to resolve the issue immediately.” Each team quickly reviewed its own monitoring data, finger-pointing began and statements like, “It’s always the network,” followed. Worse yet, the IT teams had no insight into what the branch was facing, with angry customers threatening to take their accounts to another bank down the street. Tellers were calmly trying to appease the clientele, but as time marched on, the business suffered while IT operations groups were fighting in the “war room” to find the problem.

After a couple of hours it became clear that the problem was the responsibility of the server and database teams. The server team found a hung process and the database spotted a large query that hung. Both server and database teams quickly fixed the issues and went back to their daily functions.

Unfortunately, scenarios like this one occurred often, and IT had no insight into what customers were saying and how many customers took their business elsewhere. The business suffered lost revenue and productivity.

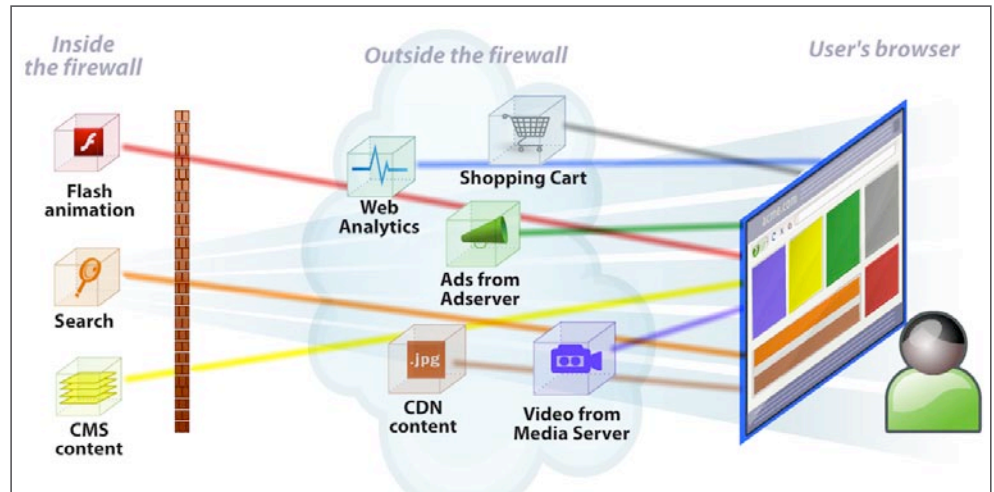
THE FACE OF THE PRESENT

As we move in time from how application performance was in the past and look to where we are today, we see there has been significant change in application architectures, IT and business ownership and responsibility, and customer expectations. The mission-critical applications are now serving many different users, both internal and external.

PRESENT APPLICATION ARCHITECTURES

The application architectures today are more complex than what they were a decade ago. What once was known as a two- or three-tiered application is now known to have “n” tiers and the application architectures are more intertwined with use of new technologies like SOA, virtualization and Web 2.0. Furthermore, applications are accessed

from multiple parts of the world, and the enterprise, which was once just in one location, is now in multiple locations that host several types of application—homegrown, Internet-based and packaged—as well as hosting the legacy systems.



From an end-user perspective, poor performance or non-availability of an application looks exactly the same irrespective of where the problem is in the service delivery chain — the service provider, the data center, the network, the enterprise, or the end user’s own device — and has exactly the same negative impact on the business.

Additionally, companies are looking to enhance these legacy systems, rather than retire them. Modernization or transformation of legacy systems allows IT organizations to keep what works and extend them to provide new user and application interfaces to the systems.

Furthermore, the applications are accessed from a variety of different devices, operating systems and browsers. Applications themselves are becoming increasingly collaborative. There is rising user demand to have these applications serve them quickly, but, more importantly, the threat of competition taking users elsewhere is even higher.

PRESENT ZONE OF RESPONSIBILITY

IT operations have a big task at hand in monitoring and managing these complex application architectures. Businesses have very high expectations of IT operations to provide well-performing applications. IT operations teams are forced to stop working in silos using their individual performance-monitoring tools and start collaborating with other silos to gather the application performance information jointly. IT operations also need to align the information with how the business users are actually experiencing application performance. This means IT operations need to start by capturing the end-user experience first, and be able to trace that back to the enterprise architecture that supports those applications.

REPORT	TYPES OF RELEVANT DATA	REPORTING LINE	VISUALIZATION
Executive	<ul style="list-style-type: none"> Service quality status Cost of poor quality Users impacted IT performance <ul style="list-style-type: none"> - % SLAs adhered to - % OLAs adhered to IT finances <ul style="list-style-type: none"> - IT revenue variance - Cost Performance Index 	CIO or Service Owners reporting to Line of Business Executive	
Service Management	<ul style="list-style-type: none"> Application downtime Service risk/SLA status User availability User downtime 	Service Manager reporting to Line of Business Executive or CFO	
IT Operations	<ul style="list-style-type: none"> Service quality of applications, transactions and locations Correlating service quality of end- user experience with key performance and availability metrics for infrastructure components 	IT managers reporting to CIO	
Troubleshooting	<ul style="list-style-type: none"> Number of slow pages Percentage of pages with large server time Percentage of users affected 	IT administrators reporting to IT management	

Role-relevant reports matter to the overall business. Being able to share the same information across the entire organization enables all parties to collaborate effectively.

However, most IT operations teams have many tools already in place and do not want to replace their performance-monitoring tool of choice with yet another one. Investing in an end-to-end application performance management solution that can take data from existing performance tools is something to consider not just for cost savings, but also for saving “ramp-up time” on learning a new tool and to leverage data that’s been collected for historical business intelligence. If IT operations continues to live in their infant state, where the silos of the operations only look at what they are separately responsible for, the business will likely suffer from loss of revenue, increased costs, less brand integrity and customer dissatisfaction.

Additionally, as organizations start to adopt new emerging technologies such as cloud computing, where the zone of responsibility is outside the

data center, parts of the application could be outsourced, and physically outside of their control; for instance, a bill-pay system. Several questions come to mind:

- Who is responsible for managing the outsourced service, the IT operations team or the Line of Business (LOB)?
- Will the data be secure if placed in the cloud?

Today, IT managers and their business counterparts need to have alignment as well as a good understanding of who owns how the service is delivered, ensuring data security and service level agreements are being met. Without this insight, the future won’t be so bright for those businesses looking to quickly grow revenue, customer base or brand recognition.

PRESENT CUSTOMER EXPECTATIONS

Today's customers are technically savvy and expect a self-serve and on-demand business experience. The need to have face-to-face interactions is becoming less a customer requirement and more a perceived nuisance. Time has become very precious to customers, so businesses must provide a means to perform business transactions on-demand via the Internet.

The business users are changing too. They tend to work less with people face-to-face and more by a means of online interactions. Businesses can place their information online for customers to access 24x7 versus a controlled store hour operation. Business users can act as operators to ensure business transactions are processed and delivered properly with little to no interaction with the customer. Businesses need to ensure application availability and quality of service because that will affect market share and revenue. If the business application isn't available or is performing poorly, customers are just one click away from using another company that will provide them a better means to make the transaction they desire.

SCENARIO 2: SURVIVAL TECHNIQUES FOR MANAGING APPLICATION PERFORMANCE

Many users and large corporations around the world bank at Paramount Funds. They perform their banking transactions online, paying mortgages or bills, transferring money in and out of IRAs and other accounts, and all sorts of other transactions. These banking transactions are done globally, from many different time zones, using many different devices and many different browsers.

On this day, the Paramount Funds branch in London has been experiencing application performance problems, and business users in Chicago are unable to access information from Tokyo. Many businesses and users, in many locations, called the support hot line complaining that they can't perform the necessary transactions. The IT operations team is put to work to resolve the issue quickly. Each IT team reviews the monitored data and alerts to seek understanding of the problem. The IT team recognizes the Internet application is experiencing problems but can't determine if it's the content served by Paramount Funds IT or other service providers.

Trying to find why the Internet application is experiencing problems became time-consuming and difficult. Even though the IT staff had powerful internal monitoring tools, they lacked visibility from the customer's perspective into the Internet application's performance and

availability. Due to the difficulty in understanding the quality of service users were experiencing, it took six days for the IT team to resolve this problem. Furthermore, Paramount Funds lost millions of dollars because the bank didn't have an integrated way to measure, monitor and identify application performance and availability across the entire application delivery chain, from the Internet and back to the enterprise.

The IT organization ultimately figured out the problem by calling one of its service providers to troubleshoot why the online banking application wasn't performing properly. After much debate, the teams from the two companies were able to fix the issue. The lack of SLAs and inability to monitor the Paramount Funds online banking application from the Internet to the enterprise was something to resolve in the future

THE FACE OF THE FUTURE

Looking forward is difficult but the future can be somewhat predicted based on current trends. Cloud computing has been hyped but gains traction as businesses see the value it brings. The promise that it offers of cost savings, scalability and agility is very attractive to many businesses, even those that plan to innovate or transform their own IT operations.

FUTURE APPLICATION ARCHITECTURES

IT will continue to extend legacy and distributed systems to add virtualization and other technologies, such as SOA, Web 2.0 or cloud computing, which will increase the complexity of application architectures. Business counterparts will continue to demand, "Do more with less," and IT will need to look for ways to save money and manage service expectations more efficiently and effectively.

Cloud computing will be the pathway to meeting many of the business needs to save money and operate more proficiently. The once-prevalent concern about "how secure will the data be if I move it to the cloud?" will continue to be a concern, regardless of who has to secure the data. But security of the application's data is now in the hands of the provider. This means IT will control less of the businesses operations, including managing data security, and become more responsible for managing service providers. IT will need to understand how to build contracts, manage agreements and track service levels over time. Additionally, IT needs a means to measure the performance of these hybrid cloud applications that are both inside and outside of the control of the business. IT will need to understand what pieces of the application are performing poorly—CDNs, ISPs, etc.

The table below is provided by the National Institute of Standards and Technology: <http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc> and explains the different models cloud computing offers.

SERVICE MODELS	
Cloud Software as a Service (SaaS)	The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based e-mail). The consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, storage or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.
Cloud Platform as a Service (PaaS)	The capability provided to the consumer is to deploy onto the cloud infrastructure applications created by the consumer or acquired, using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems or storage, but has control over the deployed applications and possibly application-hosting environment configurations.
Cloud Infrastructure as a Service (IaaS)	The capability provided to the consumer is to provision processing, storage, networks and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage and deployed applications, and possibly limited control of select networking components (e.g., host firewalls).
DEPLOYMENT MODELS	
Private cloud	The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premise or off-premise.
Community cloud	The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy and compliance considerations). It may be managed by the organizations or a third party and may exist on-premise or off-premise.
Public cloud	The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.
Hybrid cloud	The cloud infrastructure is a composite of two or more clouds (private, community or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud-bursting for load-balancing between clouds).

An end-to-end performance management solution will be needed to ensure the applications that reside both within the enterprise and outside in the cloud are monitored and measured to meet service levels and ensure the quality of service users expect. Furthermore, IT operations will need to optimize performance in the cloud application delivery chain and be able to quickly identify problems prior to users feeling any impacts.

FUTURE ZONE OF RESPONSIBILITY

As IT organizations continue to utilize cloud computing, the role of IT operations in managing infrastructure will change. IT teams will shift roles from managing and monitoring the infrastructure components, network, server, database, etc., to having a role much like a service manager who will be responsible for:

- selecting the service provider,
- defining the service level agreements
- ensuring service level compliance.

This shift will force IT and its business counterparts to be aligned and agree to the type of service delivered. Business requirements will determine what applications will be delivered through cloud computing and what IT infrastructure will remain in IT's direct control. As business owners and IT managers determine the requirements for what to push to the cloud, the role of IT will be less focused on managing infrastructure than on ensuring service providers are meeting the agreed-upon level of service. Additionally, IT Service Management frameworks such as ITILv3 or Six Sigma can assist businesses with defining processes for improving service quality and defining service level agreements.

Service Level Management					
Service Level Agreement Compliance					
Category / Service	Result	Unit	Compliance	SLA	Comments
Critical Applications					
Online Storefront	96.50	Percent Availability		Availability > 99 Percent	
Help Desk Application	100.00	Percent Availability		Availability > 99 Percent	
Critical Infrastructure					
Network	100.00	Percent Availability		Availability > 99 Percent	
VOIP	100.00	Percent Availability		Availability > 99 Percent	
Problem Management					
Sev 1 - Time to Restore	42.10	Avg Minutes to Restore		Avg < 1 Hour, Longest < 4 Hours	
Sev 2 - Time to Restore	124.20	Avg Minutes to Restore		Avg < 2 Hours, Longest < 8 Hours	
Sev 3 - Time to Restore	5.17	Avg Bus. Hours to Restore		Avg < 2 Bus. Days, Longest < 14 Bus. Days	
Help Desk					
Speed To Answer	N/A	Avg Hold Time in Seconds		Avg < 40 Seconds	
Abandon Rate	100.00	Pct Calls Answered		> 95 Percent Answered (not abandoned)	
Professional Services					
Service Project Completion	100.00	Pct Completed on Time		> 97 Percent Completed on Time	
Workstation Installations	100.00	Pct Installed in 5 Bus. Days		> 97 Percent Completed on Time	
Security Audit Completion	100.00	Pct Completed on Time		> 97 Percent Completed on Time	

Legend

Availability	All Up	Degraded	Outage
Service Level	Compliant	Below Expected	Breach

Three main objectives of measuring SLAs are to: 1. Proactively meet customer expectations and manage risk, 2. Provide input into reporting, and 3. Provide the required information for continually improving the services and the SLAs.

FUTURE CUSTOMER EXPECTATIONS

Customers will have very little tolerance for poor-performing applications, due to steep competition. The shift will continue toward self-service capabilities versus having “face-to-face” customer interactions. Additionally, users’ technical knowledge on how to work around or find something better will flourish, making companies strive to ensure they maintain high application performance to retain and increase customers.

Furthermore, technology will continue to push itself into a new era where computing is all about real-user experience. A form of hypothesis for application technology might look something like:

- Super-rich interfaces: 3-D application interfaces, hologram imaging and smart phones display pictures on any surface
- Smarter devices: easy-to-run operating systems, Internet applications for any business or personal use, quick viewing of streaming media/ video, web access everywhere
- Optimized networks and telecommunications: the new 4G vs. 3G, broadband+, vLANs, vWANs, and VPNs.

As these newer and more sophisticated ways of computing spread worldwide, customer performance demands will continue to escalate.

SCENARIO 3: THE FUTURE OF MANAGING USER EXPECTATIONS

Paramount Funds customers perform most of their transactions over the Internet using mobile and wireless devices. Many users in Hong Kong were experiencing mobile transaction problems. The online banking application allowed users to access their accounts but couldn’t perform any transactions beyond access. Businesses in the eastern region of North America had very slow performance that prevented them from performing desired transactions. Paramount Funds IT was inundated with call tickets and the IT service management team was on the job to solve the multiple regional application issues.

Paramount Funds recently implemented a private cloud, and part of the bank’s online banking infrastructure is managed by another provider. In order to resolve this issue quickly, Paramount Funds service management and the provider’s service management team continuously monitor the SLAs to determine any breach. Paramount Funds identifies a

violated SLA and works quickly to determine the source of the problem. Paramount Funds uses two CDN providers to manage serving data to users around the world. They determine that one of the CDNs is introducing delay into a critical transaction process. They quickly provide a temporary fix to redirect the users to another CDN, buying some time to troubleshoot the issue.

As Paramount Funds IT staff continues to search for the performance bottleneck, they identify the source of the problem in slow image-loading and call the CDN to communicate this finding, helping fix and resolve the performance problem.

As a result, Paramount Funds IT team manages the end-to-end application delivery chain, including outsourced services, through well-defined SLAs.

PERFORMANCE MATTERS

Achieving results that increase the company's revenues, customer satisfaction and brand perception is how to thrive in the business world today. Here are some typical results of successful APM:

When reviewing the past, present and future of application performance management, it becomes clear that application architectures, zones of responsibility and customer expectations must all be considered to manage performance and availability efficiently and effectively. As businesses around the world continue to transform and leverage new technologies, it is necessary to:

- Manage performance globally as seen by the end user
- Gain visibility into end-to-end application performance
- Align IT to meet business goals
- Optimize mean time to resolution.

Before crisis becomes catastrophe, ensure your business stays out of the news headlines about application outages and avoids performance issues that could hammer your bottom line, with application performance management that is end-to-end, from the enterprise to the Internet.

Further information on cloud computing from NIST is found at: <http://csrc.nist.gov/groups/SNS/cloud-computing/>

“Increased conversions 10%”

“Reduced homepage load time from 11.3 seconds to 3.4 seconds”

“Improved page load times 23%”

“Saved 50%+ in staff and fees”

“Reduced seven-step transaction time by 50%”

“Reduced downtime 45%”

“Achieved under 3 second response time and 99%+ availability”

“Validated decision to consolidate three data centers”

KRISTEN ALLMACHER, SOLUTION MARKETING MANAGER, COMPUWARE

Kristen has been working in the IT Service Management industry for over ten years in the areas of R&D, strategic planning, product management, sales engineering and process improvements. Kristen is a Six Sigma Green Belt and is ITIL Foundation v3 certified. Her passion is in working with people, processes and technology.

Founded in 1973, Compuware provides software, experts and best practices to ensure applications work well and deliver business value. Compuware solutions optimize application performance across the Enterprise and the Internet for leading organizations around the world, including 46 of the top 50 Fortune 500 companies and 12 of the top 20 most-visited U.S. web sites. Learn more at: compuware.com.

Compuware Corporation Corporate Headquarters
One Campus Martius
Detroit, MI 48226-5099

© 2010 Compuware Corporation

All Compuware products and services listed within are trademarks or registered trademarks of Compuware Corporation. Java and all Java-based marks are trademarks or registered trademarks of Sun Microsystems Inc. in the United States and other countries. All other company or product names are trademarks of their respective owners.

