

# Application Performance Measurement: The State of the Art

Peter Sevcik and Rebecca Wetzel

NetForecast, Inc.

Published in *Business Communications Review* February 2004

The application performance market is growing to address the burgeoning demands of networked applications and the enterprises that use them. To help managers understand the many views of performance and ensure application performance meets business needs, NetForecast recently presented a unified framework for enterprise application performance in *Business Communications Review* (BCR) [1].

In a series of articles, of which this is the first, we will use this framework to investigate aspects of the application performance market. In this article we survey tools that measure and manage the performance of information technology systems.

The market for application performance products and services can be divided into three segments: measurement, control and extension. Measurement is the art of looking into an application delivery system of servers, networks and terminals to see how well it functions and determine how to make it function better. Control vendors help deliver the best performance possible given the innate limitations of the delivery system (e.g., they maintain nominal performance when delivery conditions degrade because of congestion, loss of assets, surge of demand, etc.). And extension technologies enable vendors to further improve performance by overcoming limitations within delivery systems to bolster nominal performance to a select group of users (e.g. international users, high-value customers, business partners, etc.).

We address measurement in our initial article because, without first measuring, all efforts to optimize and enhance enterprise application performance are shots in the dark. Measurement is a cornerstone for effective enterprise application performance—and as information systems become increasingly vital to daily life, the need to measure and understand the impact of performance becomes paramount.

Just as networked applications are becoming an enterprise's lifeblood, the ways of delivering them are changing; so the ways in which their performance is measured must follow suit. This highlights the need for enterprises to reassess their measurement tools and vendors. The payoff is that

measurement and reporting provide the ability to improve performance over time.

Making sense of this complex market is a challenge, exacerbated by the fact that the market is segmented along so many axes. There are lots of vendors in the space. In researching this article, we developed a questionnaire which we sent to over 40 vendors—receiving responses from 30.

If it is a challenge for those of us who live and breathe application performance to make sense of the market, then it must be very daunting to enterprise managers. Our aim for this article, therefore, is to help enterprises understand the application performance measurement vendor landscape, and to help enterprises define appropriate criteria for evaluating vendors.

## The Changing Delivery System

In a simpler time, enterprise computer applications were exactly that—they belonged to the enterprise and were computer-centric. If you could manage performance in the datacenter and you could show that there was connectivity to your known user population, then all was well.

Critical enterprise applications no longer operate in such a simple world. Two factors govern the changing delivery system:

First, applications are no longer centralized, despite datacenter consolidation. There is a move afoot for enterprises to consolidate the application subsystems that they supply directly; however, growing numbers of those applications rely upon data sources and processes supplied by third parties.

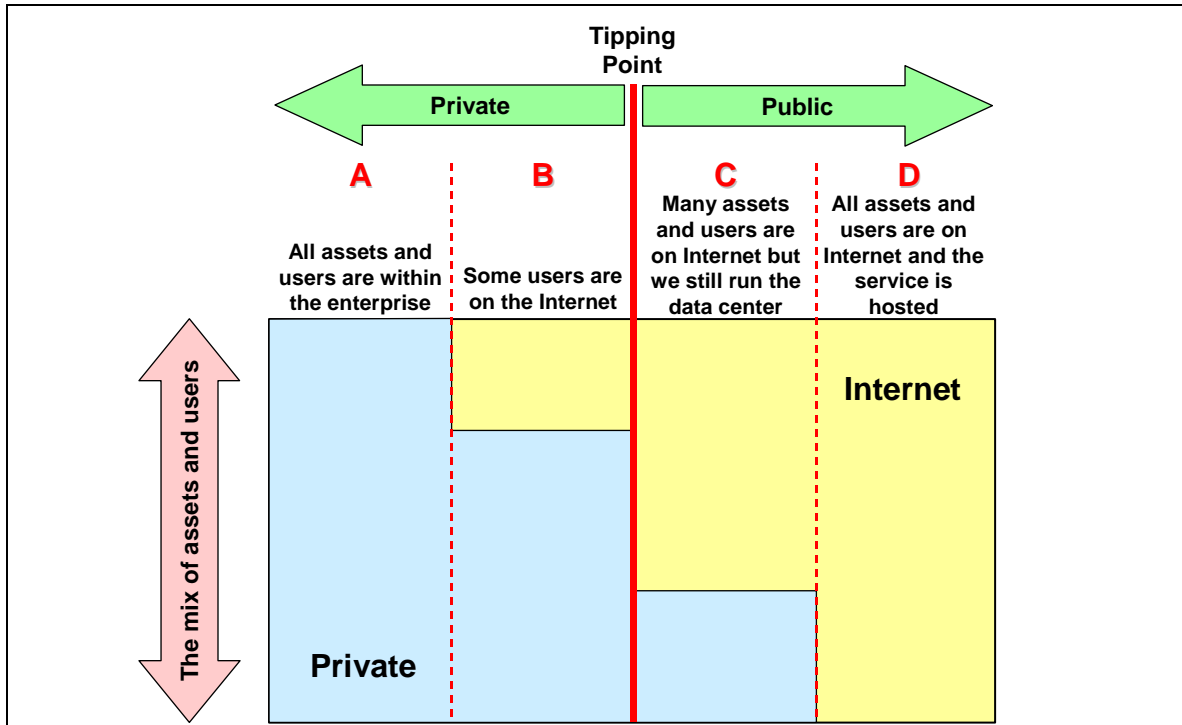
Second, today's users are often outside the realm of the enterprise—such as employees who are traveling or telecommuting, and business partners, customers, or agents. Many companies do not track who their users are and where they are at any given time. Given this situation, a measurement system that relies upon identifying and getting through to users is a non-starter.

In order to understand the landscape of application performance measurement solutions, it is important to understand the context in which these solutions

are designed to operate. This context is characterized by an irreversible migration from private to public application delivery, with assets and users under varying degrees of control by the enterprise, as shown in Figure 1. Progression from the traditional delivery system A (all assets and users encapsulated within the enterprise) to delivery system D (all assets and users on the Internet) is the natural result of demand for more efficient operations. Your enterprise moves to the

right to achieve business objectives such as improving client communication, operating a virtual company, playing to corporate strengths and leveraging the power of the Internet.

With the delivery system changing and new ways of business emerging, should you still depend upon your old private infrastructure management system? In fact, *can* it continue to measure anything meaningful?



**Figure 1 – The Transition From Private to Public Application Delivery**

**Vendor Groupings: Private versus Public**

When asked to describe what they cover, most measurement vendors have a natural tendency to answer: “We do it all.” But our research shows that the nature of their instrumentation places vendors firmly into the private/public categories defined in Figure 1.

Table 1 classifies the individual vendors based upon the Figure 1 breakdown. Many tools can be used to the left or right of their placement within Table 1. For example, an enterprise that has already

invested in a large enterprise management system (EMS) will likely continue to rely on it after it moves to scenario B. However, crossing boundaries vertically is not practical. An EMS system can’t ask for MIB variables from a public service, and a view of Internet BGP conditions is of no value inside a private network.

The following is a methodology for grouping vendors by measurement approach, as summarized in Table 1.

**Table 1 – Where the Tools are Best Applied**

		Application Delivery Scenario			
		A	B	C	D
Public-Oriented	NA	<b>Internet Service</b> Gomez Keynote	<b>Outbound</b> Adlex Coradiant Ipsium NetQoS Network Physics Packet Design Seanet	<b>Internet Service</b> Gomez Keynote	
	<b>Large EMS</b> Aprisma CA Candle Compuware Concord Fujitsu IBM Micromuse NetScout	<b>Focused EMS</b> Fidelia Quantiva Silverback Visual	<b>Data Center</b> NetIQ ProactiveNet Vieo	NA	
Private-Oriented					

Private-Oriented Measurement Approaches:

These approaches include:

*Enterprise Management Systems:* The most venerable and largest players in the enterprise application performance measurement space are the enterprise management system (EMS) vendors. Their tools use a wide variety of means to gather information about the delivery system components the enterprise owns—datacenter, remote location, network devices (switches, routers, firewalls), end user stations, etc. Enterprise management systems are large, complex, and comprehensive. A key value is their ability to scale to the requirements of a large enterprise.

Typically these tools poll devices for their SNMP MIB (Management Information Base) variables such as status and utilization from routers, switches, firewalls, and gateways (like voice over IP). Often they can also gather additional router data from Cisco’s NetFlow statistics. Some integrate RADIUS accounting data and status information from servers. Additional data may be gathered using network standards like ICMP, OSPF and the RFCs supporting ATM and frame relay transport.

A variation on this theme is the placement of RMON or proprietary probes throughout the delivery system (datacenter, LANs, WAN, remote

campuses). The probes report additional information that is not easily available from network devices, while not burdening production devices with statistics gathering.

Some solutions offer additional measurement points using active or passive agents placed in servers and on users’ desktops. Some agents are even placed in remote PCs dedicated to running synthetic transactions against the system.

A limitation of this approach is that the span of measurement stops at devices that are under the administrative control of the enterprise. These tools apply to the blue region of Figure 1. However, even though they can be used across the complete blue region, they are only cost effective in enterprises that are operating type A or B delivery systems.

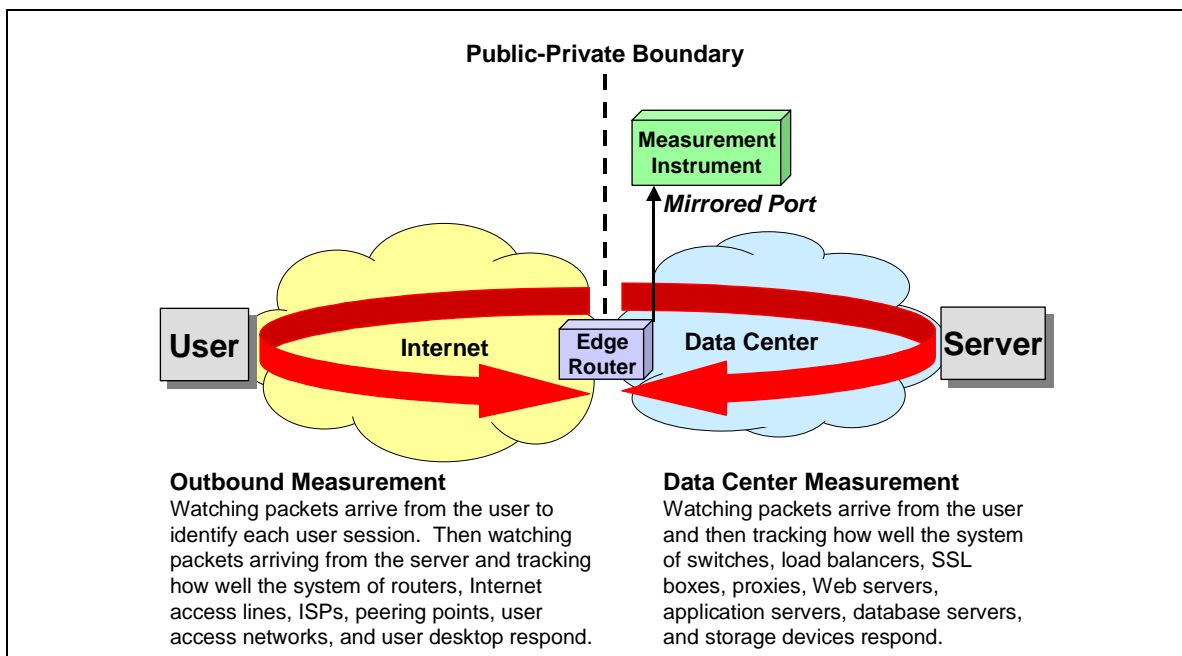
*Focused EMS:* A small group of emerging vendors is building less comprehensive, more focused solutions. Their objective is to provide simpler solutions that provide targeted data which is more relevant to specific business functions. Many of these vendors can also gather some information from the Internet to enhance their value to type B or C enterprises.

*Datacenter Measurement Tools:* Another group of emerging vendors is focused on the enterprise datacenter. These tools often sit just within the

firewall and look back toward the datacenter's network of user servers, application servers, database systems, and storage systems as shown in Figure 2. The datacenter may be the only component in the delivery system still owned by a given enterprise, but that does not make the management challenge trivial, because datacenters are growing into complex systems of machines. Understanding and managing performance at the datacenter is fundamental to any performance objective. After all, if users can't get good service at the datacenter, then they certainly can't get it over the Internet.

Although the tools in this category measure from only one strategic location, they can derive a lot of insight from the data they gather. They offer many of the features and benefits of large enterprise systems, minus the complexity and cost. Some of the tools are small versions of their EMS cousins, with probes and active agents (i.e., synthetic transactions), or passive agents.

This group of tools is best suited for the blue region of the type C delivery system. These vendors are, therefore, still grouped with the private solution vendors.



**Figure 2 – Two Views Form the Same Measurement Location**

Public-Oriented Measurement Approaches:

These approaches include:

*Outbound Measurement Tools:* All of the tools described so far are of scant value once the enterprise crosses into type C and D environments. Yes, private-oriented tools still support a portion of scenario C, but in the yellow regions, the enterprise is blind to the majority of its assets and users.

A new breed of tools looks outward towards the user from the edge of the LAN/WAN boundary [2]. These high-performance data-gathering appliances often sit off a switch port just inside the firewall as

shown in Figure 2. They derive transaction and user-specific information by reconstructing parts of the transaction from the real-time packet flow.

A tremendous amount of intelligence can be derived from these flows, which provide a detailed view from the actual users to the servers. The analysis software can calculate individual user response times (at the TCP or higher protocol layer), dynamically aggregate data in many ways, and perhaps even integrate data on the Internet based on BGP. An important aspect of this approach is that it can measure to the fine

granularity of each user without requiring a desktop agent on each of those user endpoints.

Since these tools are strategically located in a single point where all traffic must pass, they also can gather information on the performance of the datacenter and servers. The simplicity of a single appliance, combined with copious analytical power, allows some vendors in this group to provide a complete, yet inexpensive measurement platform.

*Internet Measurement Services:* The vendors in this group offer agent-based services that measure from strategic locations on the Internet, using synthetic transactions to mimic the user performing everything from a simple URL hit to a complex transaction. The views provided by these services may be physically close to the user, although they are not the actual user. Synthetic transactions may be generated from many global regions that provide good information about geographic and ISP performance. Active agents can perform standardized tests any time of the day or night, so they can test more accurately for accessibility than is possible with actual users.

#### **Assessing The Measurement Approaches**

We asked all participating vendors to place themselves within the NetForecast performance framework as described in [1]. We received 25 usable submissions that were used to assess how fully the vendor's application performance measurement supported the framework criteria. We removed the security columns (Asset Protection and User Safety) since most vendors surveyed didn't cover any of these functions, and we discuss necessary but complementary tools in "Diagnostic Tools."

We asked vendors to describe their coverage of a performance framework cell as "full" or "partial." In some cases, we made minor adjustments to vendors' submissions after follow-up communication and additional research. We believe that the results of our survey, summarized in Tables 2 and 3, accurately reflect what customers will find. Only fully-supported functions (e.g., the vendor

supplies many metrics and measurement details of the function) are granted an entry in the tables.

#### **General Observations**

As you can see, all cells are covered, with the exception of public view of the accessibility and quality of VOIP and IP video calls (see Table 3). It is possible that there truly is a lack of coverage for these functions, or it may be that we didn't happen to connect with the vendors focused on these areas.

Predictably, there are more vendors in the well-established private-oriented measurement market, but when analyzing all the entries (including partial support), the *quality* of support is equivalent in almost all cells for both the public- and private-oriented approaches.

#### **Unexpected Results**

We have preached for some time that enterprises should add views of users' experience to the more common asset views of performance [3,4]. We expected private-oriented solutions to be primarily asset managers, and public-oriented tools to focus on experience management. The results shown in Tables 2 and 3 appear to contradict this. Maybe someone is listening?

The EMS solutions have an impressive array of tools in their quivers. Some vendors offer many devices to gather unique data, and probes and agents to provide views of the user experience. However, all this coverage comes at a price. In addition to the purchase price, every new tool adds costs to operations, staffing, and complexity. It is unclear how many large EMS customers buy and use all of the comprehensive capability shown in Table 2.

By contrast, coverage of the user experience is necessarily the starting point for public-oriented systems. We were surprised by how much asset management these vendors also can provide—mostly in the form of measurement, correlation, and reporting. As a group, however, they appear weak in sophisticated problem diagnostics.

**Table 2 – Private-Oriented Measurement Coverage**

		Asset Management				Experience Management			
		Provisioning		Efficiency		Accessibility		Quality	
Real Time	Voice over IP	Aprisma Compuware Concord	Fidelia Micromuse NetIQ	Aprisma Compuware Concord Fidelia	Micromuse NetIQ NetScout	Aprisma Concord Micromuse		Concord Micromuse NetScout	
	Video Conference	Aprisma Compuware	Concord Micromuse	Aprisma Compuware	Concord NetScout	Aprisma Concord	Micromuse	Concord NetScout	
Transaction	Terminal-Host	CA Compuware		CA Compuware Fujitsu	IBM NetScout	CA Compuware IBM		CA Compuware Fujitsu	IBM Micromuse NetScout
	Client-Server	Aprisma CA Compuware	Concord IBM Micromuse	Aprisma CA Compuware Concord Fujitsu	IBM Micromuse NetScout Proactive Silverback	Aprisma CA Compuware Concord	IBM Micromuse Proactive	Aprisma CA Compuware Concord Fujitsu	IBM Micromuse NetScout Proactive
	Web	Aprisma CA Compuware Concord Fidelia	IBM Micromuse NetIQ Vieo	Aprisma Candle CA Compuware Concord Fidelia IBM	Fujitsu Micromuse NetIQ NetScout Proactive Vieo Visual	Aprisma Candle CA Compuware Concord IBM	Micromuse Proactive Quantiva Vieo Visual	Aprisma Candle CA Compuware Concord Fidelia Fujitsu	IBM Micromuse NetScout Proactive Quantiva Vieo
	Web Services	Aprisma CA Compuware	Concord IBM Micromuse	Aprisma CA Compuware	Concord Micromuse NetScout	CA Compuware Concord	IBM Micromuse Quantiva	CA Compuware Concord	Micromuse NetScout Quantiva
Data Feed	Streaming Audio	Aprisma Compuware	Concord Micromuse	Candle Compuware Concord	Micromuse NetScout Proactive	Concord Micromuse		Candle Concord	Micromuse NetScout
	Streaming Video	Aprisma Compuware	Concord Micromuse	Candle Compuware Concord	Micromuse NetScout Proactive	Concord Micromuse		Candle Concord	Micromuse NetScout
Bulk Data	Email	Aprisma CA Compuware Concord	Fidelia IBM Micromuse NetIQ	Aprisma CA Compuware Concord Fidelia Fujitsu IBM	Micromuse NetIQ NetScout Proactive Silverback Visual	Aprisma CA Compuware Concord	IBM Micromuse Proactive Visual	Aprisma CA Compuware Concord Fidelia	Fujitsu Micromuse NetScout Proactive
	File Transfer	Aprisma CA Compuware	Concord Fidelia Micromuse	Aprisma Candle CA Compuware Concord Fidelia	Fujitsu Micromuse NetScout Proactive Visual	Aprisma Candle CA Compuware	Concord Micromuse Proactive Visual	Aprisma CA Compuware Concord	Fidelia Fujitsu NetScout Proactive

Our experience shows that many enterprises do not purchase or properly integrate all of the options that are available from the complex comprehensive EMS vendors. We have often heard the line, “We have invested millions into the EMS platform but still do not have a proper understanding of performance and how to improve performance.” So, it is unlikely that many customers actually buy and receive value for their investment in experience measurement from the private group, but the obverse is not true: Customers of public-oriented tools realize the benefits of asset measurement

since it comes free and fully integrated into the products.

**Standards Are Needed**

Our survey shows many products available from an impressive number of credible vendors. But application performance measurement is complex and confusing, a problem made worse by a multitude of solutions and vendors, and this needs to be addressed by industry standards. To date, a few robust standards cover basic instrumentation, but on the analysis and reporting side, much more remains to be done.

**Table 3 – Public-Oriented Measurement Coverage**

		Asset Management				Experience Management		
		Provisioning		Efficiency		Accessibility		Quality
Real Time	Voice over IP	Ipsium Network Physics	Packet Design	NetQoS Network Physics				
	Video Conference	Ipsium Network Physics	Packet Design	NetQoS Network Physics				
Transaction	Terminal-Host	Ipsium Network Physics	Packet Design Seanet	Adlex NetQoS	Network Physics Seanet	Adlex NetQoS Seanet	Adlex NetQoS	Network Physics Seanet
	Client-Server	Ipsium Network Physics	Packet Design Seanet	Adlex NetQoS	Network Physics Seanet	Adlex NetQoS Seanet	Adlex NetQoS	Network Physics Seanet
	Web	Gomez Ipsium Keynote Network Physics	Packet Design Seanet	Adlex Coradiant NetQoS	Network Physics Seanet	Adlex Gomez Keynote Seanet	Adlex Coradiant Gomez Keynote	NetQoS Network Physics Seanet
	Web Services	Gomez Ipsium Keynote Network Physics	Packet Design Seanet	Adlex Coradiant NetQoS	Network Physics Seanet	Adlex Gomez Keynote Seanet	Adlex Coradiant Gomez Keynote	NetQoS Network Physics Seanet
Data Feed	Streaming Audio	Ipsium Keynote Network Physics	Packet Design	NetQoS Network Physics		Keynote		Keynote
	Streaming Video	Ipsium Keynote Network Physics	Packet Design	NetQoS Network Physics		Keynote		Keynote
Bulk Data	Email	Ipsium Network Physics	Packet Design Seanet	NetQoS Network Physics	Seanet	NetQoS Seanet	NetQoS Network Physics	Seanet
	File Transfer	Ipsium Keynote Network Physics	Packet Design Seanet	NetQoS Network Physics Seanet		Gomez Keynote NetQoS Seanet	Gomez Keynote NetQoS	Network Physics Seanet

The Open Group’s Application Quality Resources and Management (AQRM) working group has a solid approach to application performance measurement and management, and we encourage enterprises and vendors to join the AQRM working group ([www.opengroup.org/aquarium/](http://www.opengroup.org/aquarium/)) to make sure the needed standards are created sooner rather than later.

**Changing With The Times**

Enterprises should evaluate their current application delivery capabilities and plan for the future. Type B enterprises that ignore their “few” public users do so at their peril. IT managers who dismiss users outside the enterprise as unimportant, or believe that there is no reason to measure the public zone since nothing can be done about it, are

treading on thin ice. These old-fashioned views have their roots in type A scenarios in which IT managers measured what they could control—and they could control everything. It is fast becoming untenable for enterprises to completely own and control their IT destiny, prompting the need for new management models.

Even if an enterprise completely outsources its application delivery infrastructure in a type D scenario, it still must measure performance to effectively manage the outsourcer. Type D enterprises are likely to outsource to a type C outsourcer, which should have appropriate measurement tools for its environment. An enterprise that buys an outsourced service should

investigate how well its outsourcing vendor is measuring and managing its services.

### Diagnostic Tools

The enterprise application measurement and reporting tools described in this article provide essential, but incomplete answers for some situations, and they need to be augmented with complementary tools. These complementary performance measurement tools come into play most commonly to measure performance in pre-deployment test environments (modeling, load testing), and to diagnose problems off-line (deep analysis, what-if modeling, stress testing). Vendors of such complementary testing and diagnostic tools include: NetPredict, Opnet, Radview, Sane, and Shunra.

Most enterprises are transitioning from private to public application delivery, placing them in scenarios B or C, which implies the need for private-oriented as well as public-oriented measurement tools. In these scenarios, enterprises need integration of the results from both tool sets for best results. Although a few tools cross the public/private boundary, each provides only limited views of the other side. NetForecast predicts that this will be an area for future innovation.

Figuring out whether yours is more of a type B or type C enterprise is a complex question. (In contrast, it's simple to know if you're operating in a type A or D scenario – you control either all or none of the delivery system.) One way to distinguish between types B and C is to determine the percentage of public and private assets. Another is by percentage of users (public vs. private).

However, there are other views of this choice that may be more appropriate. For example, some enterprises run very complex Web sites that are often not operating properly—the datacenter may actually be less reliable than the Internet itself. An enterprise in this type of situation should view itself as operating a type B scenario, where monitoring of the private component is most crucial.

Enterprises need to define a process and set of business-driven objectives for this evaluation process in order to make the proper choice of a new performance measurement strategy and the vendor(s) that will support it. This decision should be made with serious deliberation based upon input from many parts of the organization and with full management involvement.

### Conclusions

Enterprise applications and needs are ever changing, and enterprise application performance measurement solutions must adapt, especially to the unstoppable migration from private-oriented to public-oriented applications. This adaptation can happen by adding public-oriented views to legacy private measurement systems, or new public-oriented tools can adapt to address both realms.

The latter is happening more quickly than the former—in a classic case of new players with new technology spotting an opportunity well ahead of entrenched players. The new players have strong offerings, which they are constantly improving. It will be a cliffhanger to see if the old players can catch up fast enough to maintain their dominance.

We predict a similar drama will play out in the control and extension segments of the performance market. We plan to survey and evaluate those areas in future articles—so stay tuned!

### References

1. Sevcik, "A Framework for Enterprise Application Performance," BCR, November 2003
2. Sevcik, "Flow-Based Management," BCR, May 2003
3. Sevcik, "Web Performance - Not a Simple Number," BCR, January 2003
4. Sevcik, "Managing Internet Performance," BCR, March 2000

### Authors

*Peter Sevcik is president of NetForecast and is a leading authority on Internet traffic, performance and technology. Peter has contributed to the design of more than 100 networks, including the Internet, and holds the patent on application response-time prediction. He can be reached at [peter@netforecast.com](mailto:peter@netforecast.com).*

*Rebecca Wetzel is an associate of NetForecast, and is an industry analyst, consultant and writer. She also is president of Wetzel Consulting LLC. She can be reached at [rebecca@netforecast.com](mailto:rebecca@netforecast.com).*



**Companies Mentioned:**

Adlex ([www.adlex.com](http://www.adlex.com))

Aprisma ([www.aprisma.com](http://www.aprisma.com))

CA ([www.ca.com](http://www.ca.com))

Candle ([www.candle.com](http://www.candle.com))

Compuware ([www.compuware.com](http://www.compuware.com))

Concord ([www.concord.com](http://www.concord.com))

Coradiant ([www.coradiant.com](http://www.coradiant.com))

Fidelia ([www.fidelia.com](http://www.fidelia.com))

Fujitsu ([www.fujitsu.com](http://www.fujitsu.com))

Gomez ([www.gomez.com](http://www.gomez.com))

IBM ([www.ibm.com](http://www.ibm.com))

Ipsum ([www.ipsumnetworks.com](http://www.ipsumnetworks.com))

Keynote ([www.keynote.com](http://www.keynote.com))

Micromuse ([www.micromuse.com](http://www.micromuse.com))

NetIQ ([www.netiq.com](http://www.netiq.com))

NetPredict ([www.netpredict.com](http://www.netpredict.com))

NetQoS ([www.netqos.com](http://www.netqos.com))

NetScout ([www.netscout.com](http://www.netscout.com))

Network Physics ([www.networkphysics.com](http://www.networkphysics.com))

Opnet ([www.opnet.com](http://www.opnet.com))

Packet Design ([www.packetdesign.com](http://www.packetdesign.com))

ProactiveNet ([www.proactivenet.com](http://www.proactivenet.com))

Quantiva ([www.quantiva.com](http://www.quantiva.com))

Radview ([www.radview.com](http://www.radview.com))

Sane ([www.sane.com](http://www.sane.com))

Seanet ([www.seanetpartners.com](http://www.seanetpartners.com))

Shunra ([www.shunra.com](http://www.shunra.com))

Silverback ([www.silverbacktech.com](http://www.silverbacktech.com))

The Open Group ([www.opengroup.org](http://www.opengroup.org))

Vieo ([www.vieo.com](http://www.vieo.com))

Visual Networks ([www.visualnetworks.com](http://www.visualnetworks.com))

NetForecast helps enterprises and vendors understand and improve the performance of networked applications. We help enterprises evaluate, improve, and manage the performance of networked business applications, and we advise vendors about customer requirements, technology issues and adoption trends.

**Smart Strategies From Hard Data**

