

Realtime
publishers

The Shortcut Guide[™] To



Assuring Website Performance through External Web Monitoring

sponsored by

neustar[™]

Dan Sullivan

Introduction to Realtime Publishers

by Don Jones, Series Editor

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We’ve made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book’s production expenses for the benefit of our readers.

Although we’ve always offered our publications to you for free, don’t think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you \$40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the “realtime” aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We’re an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I’m proud that we’ve produced so many quality books over the past years.

I want to extend an invitation to visit us at <http://nexus.realtimepublishers.com>, especially if you’ve received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you’re sure to find something that’s of interest to you—and it won’t cost you a thing. We hope you’ll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtime Publishers.....	i
Chapter 1: Understanding the Technical and Business Needs for External Web Monitoring 1	
Introduction to Assuring Web Site Performance through External Monitoring.....	2
Increasing Expectations for Service Delivery.....	2
Lessons from Early Days of Online Business	3
Lesson 1: Business Fundamentals Do Not Change.....	4
Lesson 2: Understand the User Experience.....	4
Lesson 3: Online Business Brings New Types of Risks.....	5
Lesson from Today: All Business Is Online.....	5
Key Elements of Increased Service Expectations	5
Web Accessible Services	5
Services On Demand.....	8
Consistent Service Performance	8
Implications of these Increased Expectations.....	8
Application Performance and Revenue.....	9
No Hiding Performance Problems.....	9
Increasing Application Complexity.....	10
Complexity of RIAs.....	10
Use of Multiple Back Office and Application Servers	11
Use of Multiple Partners to Provide Services	11
Increasingly Large Volumes of Data	12
Visibility into Network and Applications	12
Monitoring Server Loads and Utilization	12
Monitoring Storage Allocation.....	14
Network Performance	14
External Monitoring Meets Unmet Requirements	14
Summary	16

Copyright Statement

© 2010 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

[**Editor's Note:** This book was downloaded from Realtime Nexus—The Digital Library for IT Professionals. All leading technology books from Realtime Publishers can be found at <http://nexus.realtimepublishers.com>.]

Chapter 1: Understanding the Technical and Business Needs for External Web Monitoring

Business Web sites are the new store fronts and offices. Consumers and businesses have more options now when it comes to conducting commerce than ever before. We can travel to a retail store to browse and shop or we can go online to a retailer's Web site. Why bother calling a vendor to check on the status of our latest order when we can go online 24-hours a day, 7 days a week to see the latest updates. Rather than comparison shop for a new car by visiting multiple auto dealers, we can check inventory and prices online. For a growing number of businesses, their Web sites are their only means of contact with customers. Clearly, Web sites are now as critical to business as offices, store fronts, and call centers have been in the past.

Significant time, effort, and talent go into designing retail and professional spaces. Businesses want to provide an appealing environment for customers and clients, distinguish themselves from the competition, and strengthen their customer relationships. The same can be said for business Web sites.

In fact, there are challenges faced in Web site design that do not occur or are less difficult to manage in physical spaces. For example, navigating a Web site becomes more difficult as the volume of information increases and special consideration has to be given to ensuring ready access to different parts of the Web site. Another difference is in the level of peak demand. A retail store is limited in the number of shoppers it will have at any one point by constraints such as the physical dimensions of the store and the number of spots in their parking lot. When a customer visits an online retailer, there is no indication of how many other customers are using the Web site at the same time. (At least there should not be.) In the online realm, it is hard to hide a problem. If something goes wrong in the stock room of a retail store, most customers will be unaware of the problem. When an application server is down, a database is not performing up to expectations, or a Web server returns a page not found error, the customer knows.

Introduction to Assuring Web Site Performance through External Monitoring

Performance is a key factor in influencing the success of online commerce and other professional transactions. This guide examines the challenges facing businesses and other organizations maintaining significant Web site presence, the business drivers that create the need for adequate performance, and the technical and operational hurdles that must be overcome to effectively ensure Web site performance meets demands.

Chapter 1 examines the problem of increasing expectations for service delivery that brings with it an increased complexity of applications in an attempt to meet those expectations. This leads to the operational need for visibility into network and application performance using both internal and external monitoring.

Chapter 2 turns attention to technical issues such as the need for application-specific monitoring, reporting services, analysis and response procedures, and implementation testing. The breadth of required types of monitoring testing is clear from topics in this chapter, including Web site monitoring, application server monitoring, streaming audio and video monitoring, database monitoring, and of course network monitoring.

Chapter 3 focuses on business considerations related to Web site performance. Topics include a discussion of how business metrics are directly linked to service performance, the challenges to acquiring and retaining customers, controlling service delivery costs, and protecting brand image.

The guide concludes with Chapter 4. This chapter develops a roadmap that first builds a business case for external monitoring, second, outlines a plan for a monitoring framework, and finally shows how to design and implement an external monitoring program. Testing and acceptance criteria as well as long-term maintenance and governance issues are discussed as well.

Let's begin with a look at how expectations for service delivery have changed with the advent of online commerce.

Increasing Expectations for Service Delivery

Conducting business online can be fast, cost effective, and profitable. The evolution of commerce has accelerated with the advent of the Internet and in the process created a trend of increasing expectations. Today, it is not considered unreasonable to conduct business at any time, from anywhere, and with any business. Moving service delivery online was not without its difficulties, as the timeline in Figure 1.1 shows.

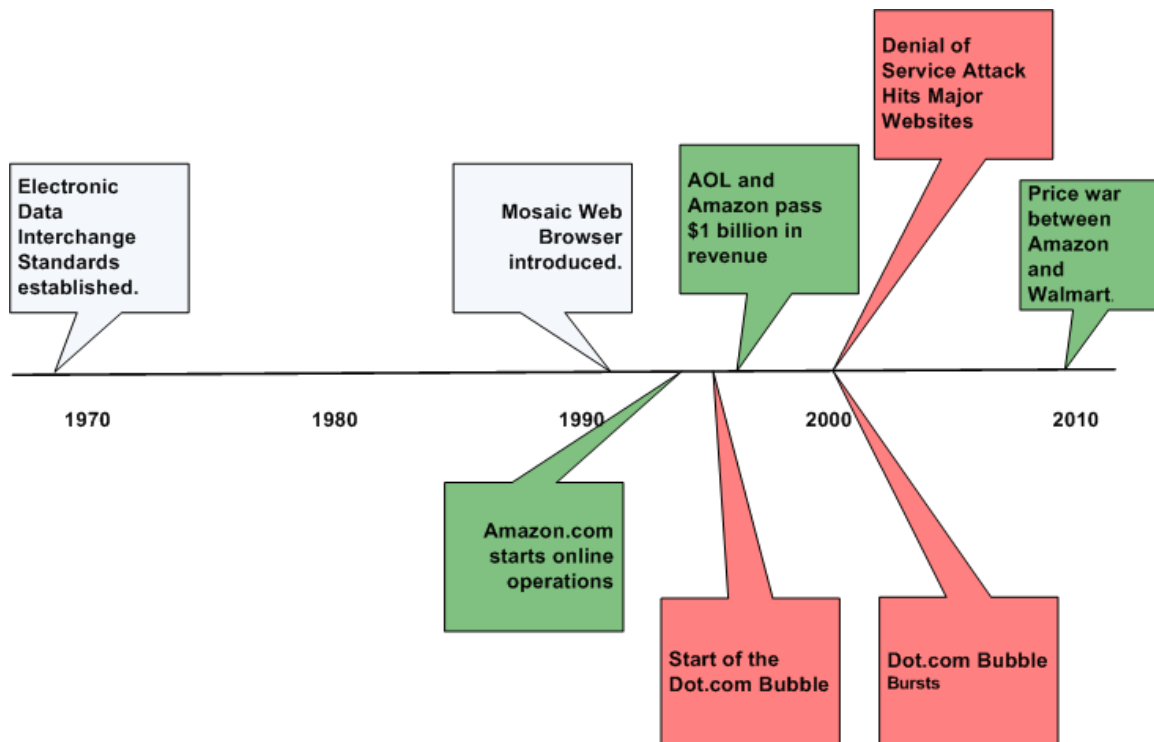


Figure 1.1: Online business has its roots in early data exchange standards of the 1960s but accelerated after the introduction of the Web browser, resulting in spectacular successes and failures as well as the eventual blending of online and brick-and-mortar-based business.

Long before the Internet was adopted for commerce, businesses were making use of electronic data exchange. The Electronic Data Interchange (EDI) standards were adopted in 1968 and provided the foundation for basic, structured information sharing between message trading partners. EDI provided an early example of how the cost and time required to conduct a business transaction could be significantly reduced using information technologies.

Lessons from Early Days of Online Business

Electronic commerce remained largely a back office matter until the 1990s with the introduction of the Web. Almost 25 years after the introduction of EDI came the invention of the Web browser as we know it today. At that point, ecommerce began to take off, and it would not be another 25 years before the next significant milestone. Within 3 years of the introduction of the Mosaic browser, Amazon.com began conducting business online. By 1998, the online bookseller had broken \$1 billion in sales with that year's holiday shopping season. They were not alone. By 1998, the dot-com bubble was in full swing, and within 2 years, the bubble had burst.

Keeping in mind the well-known warning that “those who cannot remember the past are condemned to repeat it,” it is worth noting some of the lessons learned from the dot-com bust. They can be summarized as follows:

- The fundamentals of business do not change by moving transactions online
- The customer’s online experience is strongly tied to long-term success
- New risks come with online business

Let’s consider some examples of each of these lessons.

Lesson 1: Business Fundamentals Do Not Change

Conducting business online seemed to be the wave of the future, and there were plenty of ideas on how to capitalize on that. Some succeeded and many failed. Those that failed show the perils of online commerce can come from traditional business concerns. For example, you have to earn a profit no matter what you sell or how you sell it. There was no “new set of rules” for ecommerce; the fundamentals of business have not changed since the ancient sea-faring Phoenicians started trading with their Mediterranean neighbors. There were however, new technical challenges to overcome as businesses learned to adapt Web technologies to business.

Pets.com is a classic example of a company with an unsustainable business strategy. Known for its clever advertising, Pets.com coupled a sock puppet mascot with the idea of selling pet food over the Web. Unfortunately, the shipping costs for a 40-pound bag of dog food are much higher than the cost of shipping a couple of paperbacks from an Amazon distribution center. Couple that with the fact that although we can often wait a couple of days for a book to arrive in the mail, hungry pets are less patient. To attract customers, the company subsidized shipping costs to the point of losing money on many transactions. Pets.com began operations in 1998 and ceased operation in 2000.

Lesson 2: Understand the User Experience

The online fashion retailer, Boo.com, took on the complex challenge of becoming a global retailer in 1998. The company was based in the United Kingdom but had facilities in other European countries as well. The company had plenty of business hurdles to overcome including navigating multiple tax jurisdictions and converting between many currencies (the euro was not yet fully adopted in the European Union at that point). Even if they could manage the strictly business issues, they faced technical limitations that could have undermined the entire enterprise. The complex Boo.com site did not perform well for many customers who were limited to dial-up connections. It was like a customer visiting a brick-and-mortar store only to find she cannot get far past the front door. As Tritan Louis, former interim CTO of Boo.com, put it, “No matter how good your backend systems are, the users will only remember your front end. Fail there and you will fail, period.” (Source: Tritan Louis, “Boo.com Goes Bust “at <http://www.tnl.net/blog/2000/05/19/boocom-goes-bust>.)

Lesson 3: Online Business Brings New Types of Risks

In February 2000, a number of major Web sites became victims of a Denial of Service (DoS) attack. Yahoo.com, Amazon.com, Buy.com, eBay.com, and others were subjected to a flood of Internet traffic designed to disrupt services. The attack showed how the technologies that enable online business can be turned against it. It also made it clear that consistent performance cannot be taken for granted on the Web. Networks can become congested, servers may be overloaded, and applications can fail.

Lesson from Today: All Business Is Online

History highlights important facts about online business, but perhaps the most important is demonstrated today: The lines between online business and other means of commerce are blurring. A recent *New York Times* headline captures this with “Price War Brews Between Amazon and Wal-Mart” (Source: *New York Times*, November 23, 2009, at <http://www.nytimes.com/2009/11/24/business/24shop.html>). The seeming unassailable retailer known for its brick-and-mortar operations is being challenged as much, or more, by an online retailer than by traditional competitors. We are close to the point where the distinction between online and offline business may no longer be a useful concept when formulating business strategy. To understand why this is happening and how to compete in this new environment, it is important to understand the key elements of increased service expectations.

Key Elements of Increased Service Expectations

We have come to expect more of business services with regard to online services; in particular, we expect:

- Services to be Web accessible
- Services to be available on demand
- Consistent performance from those services

Meeting each of these expectations requires a complex structure of multiple technologies and business practices.

Web Accessible Services

In its most basic form, a Web-accessible service is available through the use of standard Internet protocols. A file transfer site where customers or business partners can transfer files using the File Transfer Protocol (FTP) would constitute a Web accessible service by this definition. This type of basic functionality is important and is often included in more complex services, but it does not begin to capture the full range of services that can be provided by the Web. When you think of Web accessible services, it helps to think of them both in terms of business functions, such as placing an order, and in technical terms, such as creating an XML structure and passing it to a Web service. Before you can deliver any type of service, you have to understand the features and the constraints of the delivery platform.




Figure 1.2: Web accessible implies accessible from multiple types of devices and operating systems.

Supporting Multiple Platforms

We can access the Web from a wide variety of platforms. When the client/server model of applications developed, we had to ensure our applications would run on the different operating systems (OSs) that might be found on a local area network (LAN) or corporate wide area network (WAN). With Web applications and common browsers across OSs, there are fewer barriers to delivering applications to Windows, Mac, and Linux desktops and laptops. Today's cross-platform challenges center on the needs of mobile devices, especially those with small form factors, like smartphones.

Supporting Multiple Business Functions

We are increasingly expecting to carry out any relevant transaction with a vendor online. For example, if we order a product online, we expect to be able to track its shipping status or cancel the order. When we bank online, we expect to be able to transfer funds between accounts, make payments, and check balances. Many of us have become accustomed to carrying out internal business functions online as well, such as checking the status of HR reports or submitting internal forms online.

Supporting Multiple Backend Technologies

There are many components within IT infrastructure that are required to make services Web accessible. Web servers, application servers, databases, directory servers, authentication and authorization mechanisms, and network monitoring are just the starting components for some of today's complex information processing systems. In some cases, a single vendor may be able to provide a full suite of applications, but often, distributed systems are made up of components from different vendors, using different development platforms (for example, Java and .NET), and running on different OSs. Service-oriented architectures help hide implementation details and allow developers and architects to work the common level of abstraction of Web services. There is, however, significant complexity in the application stack required to make back office applications Web accessible.

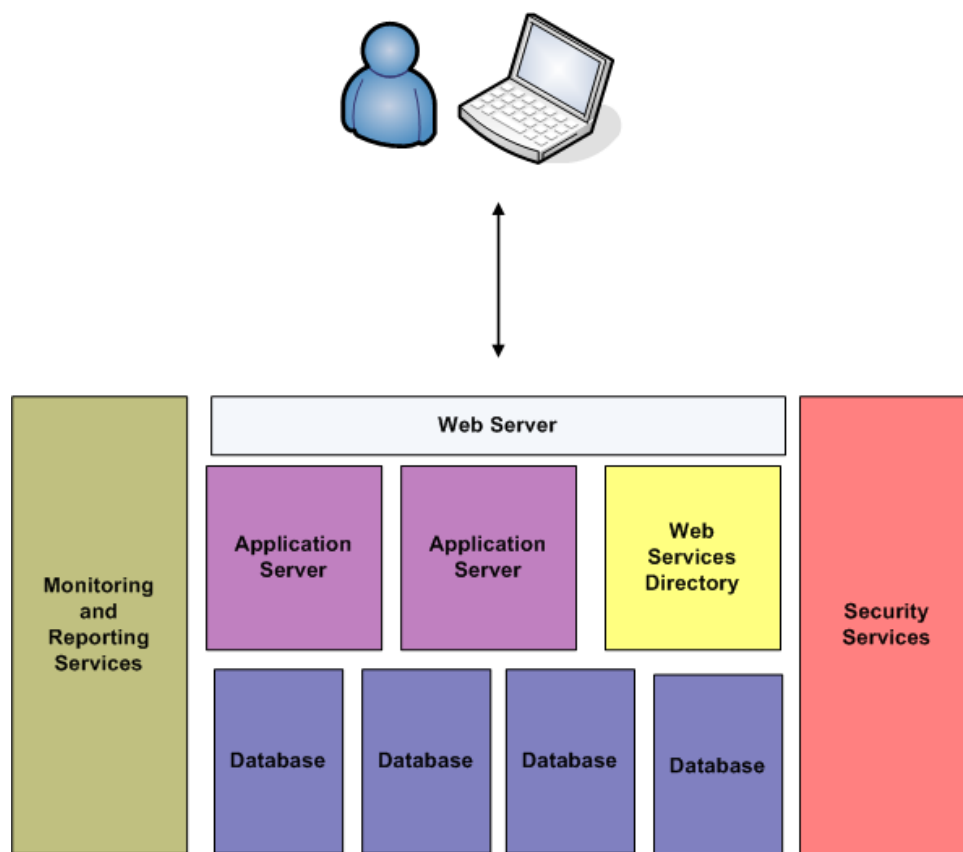


Figure 1.3: Many types of services are required to create and maintain Web-accessible applications, most of which are not visible to application users.

Services On Demand

In addition to having a wide array of services available to us, we also expect these services to be available on demand. We do not expect stores, banks, or business offices to be open around the clock, but businesses have extended their ability to work with their customers after hours. For example, with the proliferation of call centers, many of us have become accustomed to calling businesses outside of normal business hours. Similarly, we expect those businesses to maintain continuously available Web sites. The expectation of services on demand is closely coupled with our expectations for accessing services through the Web. The Web allows us to work and conduct business from anywhere we have basic connectivity and because the Web “never sleeps,” why should the services that are made available through the Web?

Consistent Service Performance

Implied in the expectation for Web accessibility and on demand services, we expect consistent performance when we use these services. In addition to the architectural complexity of integrating multiple components to deliver Web accessible services, you now have to address the need to keep these systems running constantly. This introduces additional challenges:

- Scheduling maintenance on servers without disrupting services
- Replicating data from primary to standby databases in the event of a failure on the primary server
- Rapidly detecting component failures or degraded performance
- Scaling hardware and applications for peak demand while maintaining reasonable utilization rates

The issues of Web accessibility, service on demand, and consistent performance are primarily technical issues that must be addressed by developers, designers, and architects; however, the implications of these expectations extend to other areas of the business.

Implications of these Increased Expectations

Two facts have become clear about delivering services over the Web:

- Application performance has a direct impact on revenues
- There is no hiding performance issues on the Web

There are comparable problems in traditional commerce, but as you shall see, they are particularly problematic for online service delivery.

Application Performance and Revenue

Earlier in this chapter, a CTO of a failed dot-com company made the point very clear that if the application interface hinders the user experience, you lose the customer, period. There are unfortunately, many ways for performance to lag behind customers' expectations:

- Servers not keeping up with growing demand, especially during peak periods, such as the holiday shopping season for retailers
- Insufficient network bandwidth to key parts of a distributed system; average speeds are irrelevant; application speed is limited by the weakest link in the chain
- Inefficient interface code that requires significant data transfers between servers and client devices
- Difficult to navigate Web sites that require too many steps to complete a transaction leading to high abandonment rates

Each of these problems can adversely affect revenue. Execution is often a distinguishing factor among competitors, but it is especially important in online service. A customer who might have to drive to a competitor's store might hesitate to leave but that same customer might not think twice when the competitor is just a search term and a few clicks away.

No Hiding Performance Problems

As long as a business can keep the customer-facing offices, stores, and facilities well staffed and reasonably functional, customers might never suspect any dysfunction in the back office. That is not the case with the Web. When a server is down, customers will know it (unless there is a high-availability backup server with automatic failover in place). You might be able to gloss over the system failure with a reassuring error message about how you value the customer's business and promising to work diligently to resolve the problem, but in the end the customer knows you failed.

Customer reviews on opinion sites add another dimension to the inability to mask performance problems. If problems are severe enough and affect enough customers, comments will turn up on any number of the popular customer opinion sites. Imagine customers complaining about the inability to complete an order, difficulty navigating a site, or long delays in getting an order conformation. If that happens, the impact of the failure may be felt long after the problem is solved. Increasing expectations around service delivery is only one aspect driving the need for external monitoring of Web service delivery; the increasing complexity of applications is another.

Increasing Application Complexity

Web applications have come a long way from basic HTML pages with simple client inputs—a single server responds. Web applications are more like desktop and client/server applications with complex interfaces, multiple backend servers, and significant volumes of data. Each of the following aspects of this increasing complexity should be considered when planning for operational testing and long-term maintenance:

- Complexity of rich Internet applications (RIAs)
- Use of multiple back office and application servers
- Use of multiple partners to provide services
- Increasingly large volumes of data

Any one of these factors would add to the difficulty of designing and maintaining Web-based services; the four together compound the effect.

Complexity of RIAs

RIAs get their name from the wealth of features that are incorporated into Web applications. RIAs are often compared with desktop applications and provide advanced functionality by incorporating developer frameworks, such as AJAX, Adobe Flash, and Microsoft Silverlight. Frameworks such as these enable:

- Client-side processing so that not all requests have to be sent to the Web server, making the interface more responsive
- Reduced data exchange, decreasing demand on network resources
- Application state information maintenance, enabling more complex processing on the client side

This advanced design model takes advantage of the processing power and memory of client devices to reduce demand on the network and servers while improving the functionality of the user interface. It is important to understand, though, with RIAs, significant amounts of application functionality are isolated to the client device, which is often outside the control of the Web application provider. Thorough testing during development is required, along with ongoing monitoring during operational use, to ensure the RIA is functioning properly and meeting requirements.

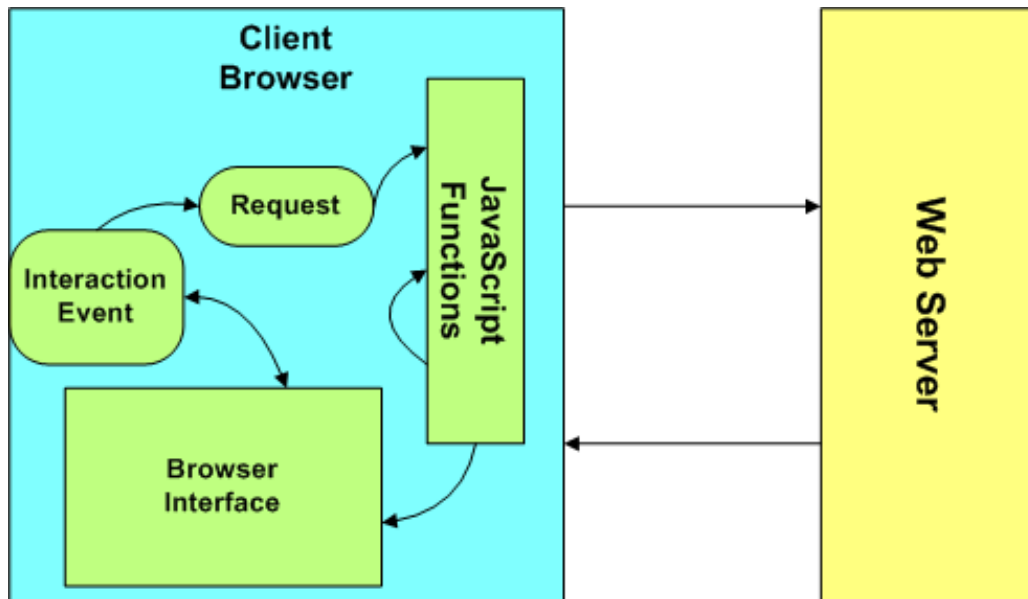


Figure 1.4: RIAs enable more processing on the client side than traditional HTML interfaces allow, increasing the complexity of testing and monitoring required to ensure proper and consistent functionality.

Use of Multiple Back Office and Application Servers

It may not be apparent to a user, but there are often multiple servers running behind the scenes of today's business applications. Requests from clients may go directly to a single Web server, and that request triggers a series of requests to other servers. For example, a Web server may retrieve content from a content management server, which in turn makes a request to an authentication and authorization server that depends on a relational database for low-level storage services.

Some applications make use of more independent sets of servers. In the case of mashup applications, content from multiple providers is aggregated at the interface level rather than at the application level. An example of a mashup would be a stock portfolio analysis application that displays summary information about a customer's portfolio along with graphs retrieved from a third-party service, news culled from business sources, and current stock market summaries from yet another service. This type of application exemplifies another common technical challenge: services provided by multiple partners.

Use of Multiple Partners to Provide Services

Complex business services require a specialization of services. A retailer, for example, may excel at marketing and inventory management but lack the legal expertise to maintain sales tax calculations on a national or global scale. The retailer also depends on shipping services provided by third parties. From the customers' perspective, we are doing business with a single vendor not three, and if there is a problem with the tax on an order or the speed at which it is delivered, it reflects on the retailer, not necessarily their business partners.

When monitoring complex applications, you need to understand dependencies on third partners. If a partner's services are on a critical path—for example, a sale cannot be completed without calculating taxes—then that service as well as equally important internal services should be monitored.

Increasingly Large Volumes of Data

Applications give customers access to growing volumes of data, in the form of catalogs, documents, and audio and streaming video. Ten years ago, a home improvement store may have made their catalog of products available online with descriptions, pricing, and availability information. Today, that same business can offer instructional videos for the do-it-yourself customer, links to manuals and other product documentation, and detailed information about order status, inventory levels, and product availability. This increasing volume of data converges in the end user application after flowing through a complex web of applications, servers, and network routes.

The combination of RIAs at the front end and multiple servers and multiple partners on the backend along with increasing volumes of data make it more difficult to predict and manage application performance. With this type of complexity comes the need for information about multiple points in the overall application architecture.

Visibility into Network and Applications

As we have seen in this chapter, there are many components in today's Web applications that, along with increasing expectations about performance and availability, require you to have frequent insight into the full range of subsystems. In particular, you should monitor:

- Server loads and utilizations
- Storage allocation and use
- Network performance

These are areas that have required monitoring as long as there have been distributed applications, but the process is becoming more difficult.

Monitoring Server Loads and Utilization

Monitoring servers sounds like a pretty straightforward requirement and should be reasonably easy to implement, at least in theory. After all, even basic monitoring tools such the one that Figure 1.5 shows can give a picture of the overall load on a server. Performance monitors can help pinpoint bottlenecks in memory, disk, and network use and show variation in CPU load over time. If the entire application ran on one server, a basic performance monitor might be enough to understand server loads. That is not the case.

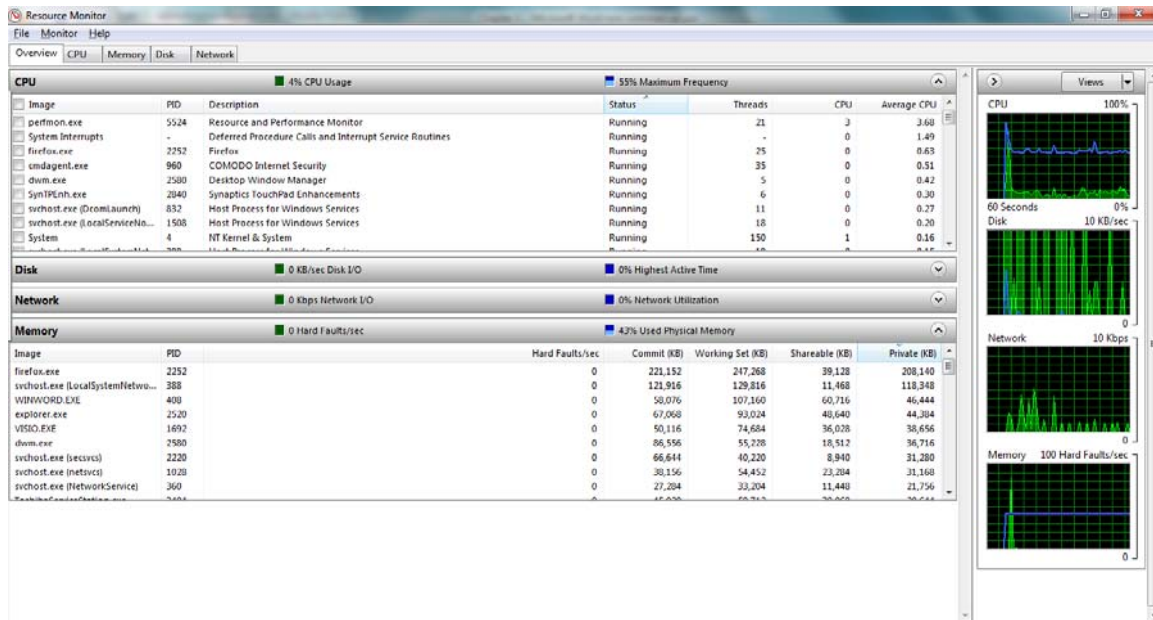


Figure 1.5: Monitoring individual servers with local monitoring tools is just the start to the type of comprehensive monitoring that is required to ensure effective delivery of online services.

Web applications are often distributed over multiple servers by their nature. Performance monitoring has to occur over all the servers, including Web servers, application servers, database servers, proxy servers for back office mainframe applications, and so on. In addition, you need to be able to correlate events across servers. For example, do peak demands in application servers correlate with peak demands in database queries? What is the effect of increased write transactions on the network? Is there a bottleneck, such as with the authentication and authorization service, that is slowing other application response times?

Under-utilization is also important to catch. When server loads are lower than necessary, costs may be reduced by combining servers, often using virtual machines, on a single physical server. By monitoring servers over extended periods of time and under different conditions, you can understand the baseline utilization rates and more efficiently allocate physical resources.

Monitoring Storage Allocation

Storage use can rise and fall with application use. You should understand your storage requirements in terms of baseline needs—that is, storage requirements under common usage patterns—as well as in terms of trends. Ideally, there would be an easily discernable relation between application use and storage requirement; for example, if the number of users doubles, so does the storage. This is often not the case. There are fixed storage requirements that exist regardless of how the application is used, there are variable storage requirements that change with application use, and there are storage demands driven by other business requirements, such as the need for high-availability failover services. Storage allocation monitors can collect the information needed to make both short-term management and long-term architectural decisions.

Network Performance

The best configuration of servers running optimally-designed software will fail from a user's perspective if the network performance is poor. A network is a unified collection of services and hardware that can operate effectively under a wide range of conditions, but it should not be taken for granted. In particular, application engineers and systems administrators should have access to data about:

- Performance statistics, such as bandwidth utilization, packet loss, and latency
- Device-specific data, such as loss rates and packets to and from the device
- Log information collected from devices
- Simple Network Management Protocol (SNMP) data

Even with comprehensive monitoring of these aspects of the application environment within a corporate network, you can still miss critical information. Internal monitoring of servers, storage, and networks is necessary but not sufficient to collect all the data needed to understand application functionality from a customer's perspective. In fact, there is no single customer perspective. A customer in one geographical region may experience significantly different performance than those in another region. Two users in the same area but using different client platforms can also have widely different experiences. The only way to capture this external data is by external monitoring.

External Monitoring Meets Unmet Requirements

Think of monitoring in terms of expanding scope. You should choose a scope appropriate for the type of application you are monitoring (see Figure 1.6). For example, for applications that run on a local device, such as a desktop application or a locally installed graphics package, you would employ local monitoring. A client/server application that makes use of locally connected desktops and servers requires internal monitoring of servers, client devices, and LAN performance. As you move outside the local network, any online services you provide to customers outside your firewall will require external monitoring to ensure performance. These will include your simple informational Web site pages, transactions such as logins and purchases, and streaming video and audio as well as any services you provide through a third party.



Figure 1.6: Different types of monitoring provide varying level of details with perspectives that range from device-specific to a global application experience.

External monitoring is the process of collecting data about application performance from outside the corporate network. External monitoring complements internal monitoring by capturing data about end customer experience. For example, customers in the United States may have a significantly different response time than customers in Australia; in fact, customers in New York City may have a different response time than customers in Dallas. Only by external monitoring can you identify problem areas outside corporate control and collect data that will help you resolve or at least avoid problem areas.

Delivering consistent and acceptable execution is essential to maintain the trust of customers. With external monitoring, you may be able to detect early warnings of potential problems before they become significant threats to performance. In the case of business to business transactions, external monitoring can document the fact that service level agreements (SLAs) are being met. External monitoring should include detailed reporting for both performance analysis and SLA reporting. Reports should target information to monitor key performance indicators, such as response time by core application function, as well as alert systems administrators to errors or potentially problematic performance issues.

Summary

You need a broad array of monitoring tools to help ensure Web applications perform as needed. Customers increasingly expect consistent performance, high availability, and ubiquitous access to information and services. Meeting these expectations often entails building increasingly complex applications built on increasingly complicated architectures. Today's Web applications are complex combinations of services provided by different technologies, running on different servers, managed by different business partners. It is hard to imagine how you could have created more complexity in IT if you tried. No matter how well designed the system, once it is deployed and in use, it will require constant monitoring. Web applications have created additional types of monitoring requirements, especially with regard to understanding performance from a user's perspective. External monitoring can meet the challenges created by increasing expectations and complex applications. The next chapter will focus on essential components of external monitoring.

Download Additional Books from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this book to be informative, we encourage you to download more of our industry-leading technology books and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.