

Realtime  
publishers

# *The Definitive Guide™ To*

# Monitoring the Data Center, Virtual Environments, and the Cloud

*sponsored by*

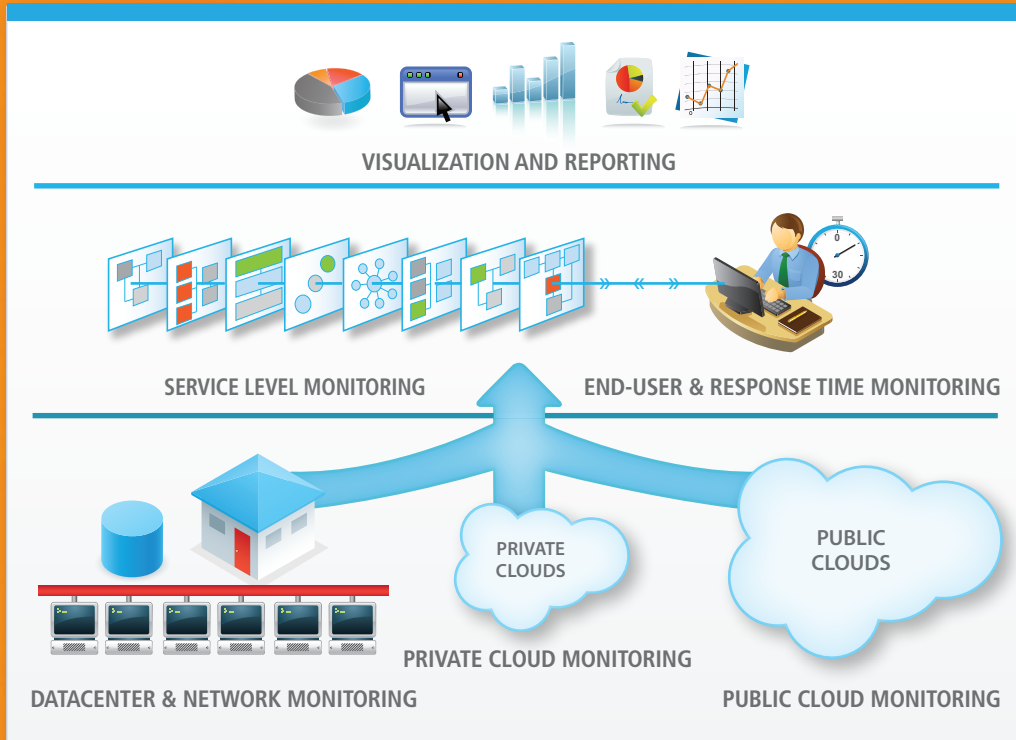
nimsoft

From the Datacenter to the Cloud

*Don Jones*

# The Nimsoft Monitoring Solution

Unified Monitoring



- Ensures business service delivery regardless of IT platform
- Enables rapid adoption of new computer infrastructure such as private and public cloud
- Monitors the datacenter to the cloud, including SaaS, hosted, and virtualized environments
- Lowers TCO by 80% and delivers proven value in weeks

Chapter 6: IT Health: Management Reporting as a Service.....	88
The Value of Management Reporting .....	88
Business Value .....	88
Technology Value .....	89
Reporting Elements .....	89
Performance Reports .....	89
SLA Reports.....	106
Dashboards.....	108
The Provider Perspective: Reports for Your Customers .....	110
Conclusion .....	111

## **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 or 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 email at [info@realtimepublishers.com](mailto: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 6: IT Health: Management Reporting as a Service

---

I've spent a lot of time in this book explaining the capabilities and technologies you need to add to your environment in order to enable truly hybrid, data center-to-the-cloud application and service monitoring. But all of that monitoring is useless without output: One of the end goals of this entire effort is to provide your managers and executives with effective reports—whether they are internal “customers” or external customers. Dashboards and other elements that show a manager that the environment is healthy and on budget or that show them which *service* (not IT component) isn't doing well. The goal of this final chapter is to focus on these reports, what they should look like, and what value you can expect to derive from them.

### Note

This is an unusual chapter in that I'll mainly be presenting examples of reports. My goal is to help you develop a kind of “shopping list” for the types of reports you should look for in systems that you're evaluating and to explain some of the finer details that I like in these reports. Most of these examples are taken from live systems, so in some cases, I've obfuscated customer-specific information such as publicly-accessible server names, IP addresses, and so forth.

### The Value of Management Reporting

There's no question that reporting has value, but what, specifically, is that value? In other words, what should you expect reports to provide other than pretty graphs? What will you *get* out of reports? Let me quickly outline the major points so that I can then show you examples of monitoring system reports that deliver those benefits.

#### Business Value

Businesses look, primarily, for reliability and return on investment (ROI). Specifically, businesses want reports that can:

- Monitor compliance with service level agreements (SLAs)
- Monitor application performance from an end-user perspective
- Track utilization, especially when that utilization relates to cost, as it does with most cloud-computing platforms
- Help predict growth in utilization (to help estimate the costs of supporting that growth)
- Assist in maintaining maximum uptime and responsiveness for entire applications

## Technology Value

Technologists need reports and tools that can help them achieve the business' goals. That means technology-focused reports are often a bit (or a lot) more detailed, focusing on implementation details that support the business' high-level views and metrics.

Technologists look for reports that can:

- Quickly detail key performance metrics for components, highlighting out-of-tolerance metrics that require attention
- Show usage trends so that IT can predict when usage will exceed the system's ability to perform within tolerance
- Dive from high-level metrics, like end-user experience measurements, into deeper, technology-specific metrics for troubleshooting purposes

## Reporting Elements

In the examples that follow, I will highlight specific capabilities of a monitoring system. In most cases, I'll call out specific features of these reports that I find especially useful, and that I think you should look for in your own monitoring solution. I'll spend the most time on detailed performance reports because those provide the bulk of the intelligence you'll need to operate your infrastructure. I'll also look at SLA-specific reports and a few dashboards that help provide a high-level, at-a-glance view of the environment or specific applications and services.

## Performance Reports

Let's dive into the examples I've gathered. First up, in Figure 6.1 is a look at Exchange Server availability. This report really highlights the value of being able to monitor a hybrid IT environment: *This* Exchange system is hosted at Rackspace, not in our own data center. Being able to monitor system uptime—especially when other applications depend on this system—is crucial to maintaining the overall performance of our environment.


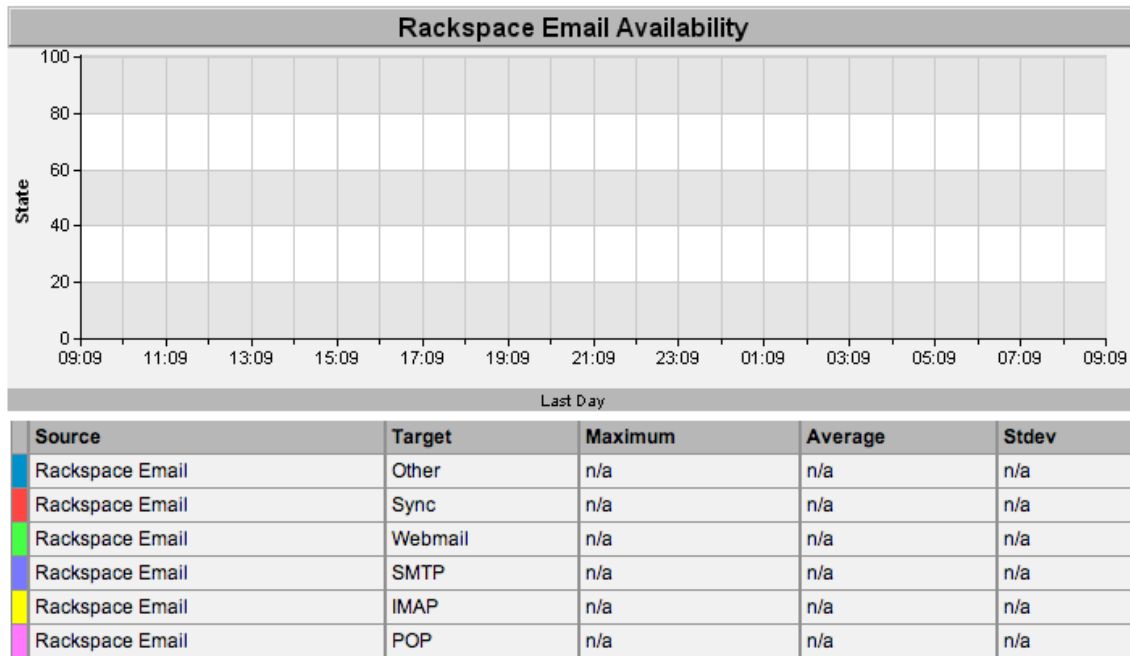


Figure 6.1: An example Exchange availability report.

**Note**

You're seeing a really good day on my Exchange Server; here the performance line is the blue one at the very top of the graph, indicating 100% uptime.

Figure 6.2 shows another way of monitoring email availability, and it illustrates a key capability that you should look for. This chart is showing overall email availability from a variety of services—Web mail, SMTP, POP, and so on. Those are the *services* you rely on, so rather than monitoring the *system*, this report is monitoring those services. This kind of service-level availability is important for anyone who is relying on hosted or cloud-based services as a part of their IT infrastructure.



**Figure 6.2: An example general email availability report—rackspace hosted services.**

**Note**

Once again, everything's looking good—all of these services are at 100%. Boring-looking performance charts are the ones you hope to see all the time!

Figure 6.3 shows another hosted element: Salesforce.com statistics. This is a more basic performance report, showing the number of transactions as well as a couple of service level-type statistics: transaction speed and overall system status. There are numerous other stats you would want to track for Salesforce.com if you relied upon it, and your monitoring system should deliver in this kind of easy-to-read, *live* report.

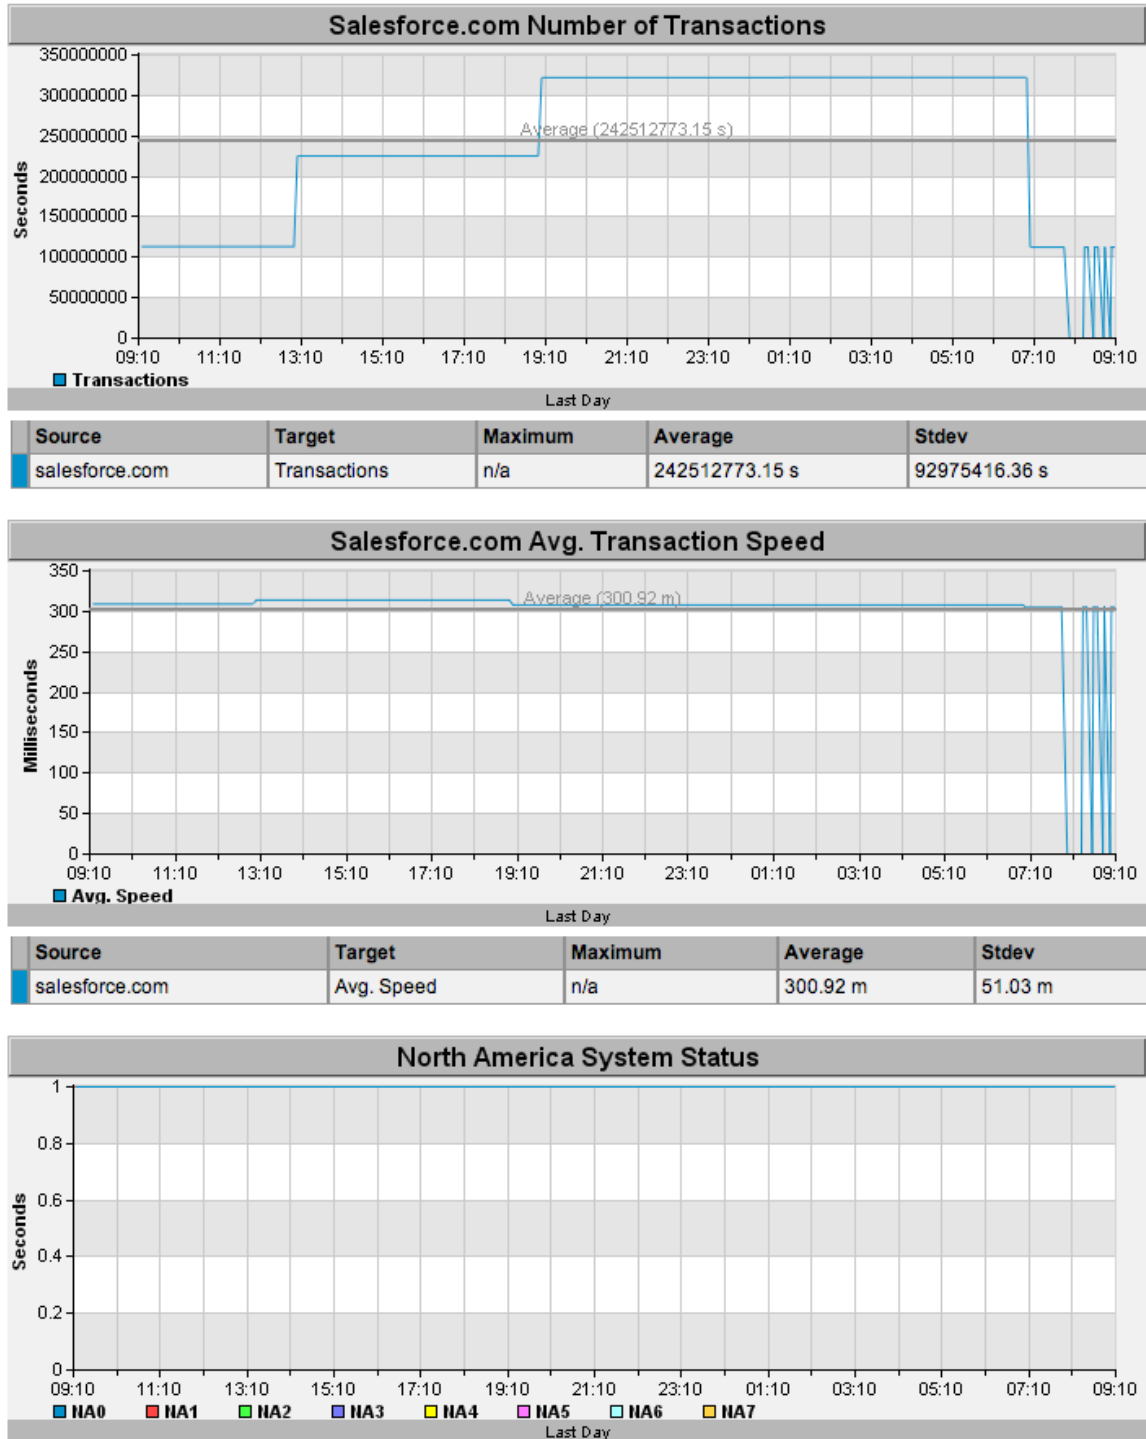


Figure 6.3: Example Salesforce.com statistics.

**Note**

Notice that a drop in the number of transactions isn't bad, although the drop in transaction speed might be worrying. Neither of these statistics affects the system's total uptime, shown on the bottom graph as 100%.



If you're using a network (and what else would you be using?), you should be concerned about its performance and availability. Figure 6.4 highlights another key capability I've discussed throughout this book: getting everything into *one* monitoring system. Just because you *can* monitor network protocol statistics using other tools, you should still want them monitored in the same place as everything else. That way, when a problem occurs, you have all the troubleshooting information you need in *one* place.


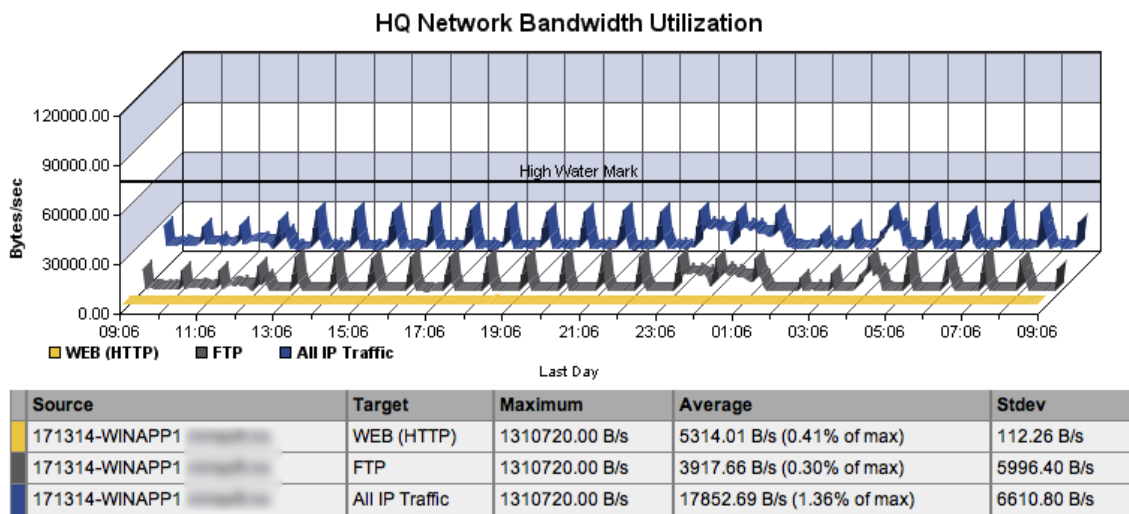


Figure 6.4: Network statistics from Cisco Netflow.

**Note**


No DNS traffic at all? That could potentially be a bad sign, except that in my network, most DNS is resolved internally, so we're not seeing DNS use much bandwidth.

Figure 6.5 is another example of a service-level report, showing overall network bandwidth utilization. I especially like the inclusion of a “high water mark” line, showing where bandwidth has maxed out in the past 12 hours. Notice that there's also a break down of protocol traffic, so if there *is* a problem, you can get a good idea of what protocol is contributing to that problem.



**Figure 6.5: Network bandwidth.**

If you have any service or application that relies on external services—such as an external LDAP server—then you need to be able to monitor that. Figure 6.6 shows how a monitoring system can do so, connecting to an external LDAP system and measuring response times. By establishing health thresholds for these response times, you can start to create alerts and other notifications when response times exceed your tolerances.



**Figure 6.6: Service response time—LDAP.**

Traditional server monitoring should be included as well, as illustrated in Figure 6.7, which shows common stats for a Windows server. Again, it’s not so much that you don’t already have monitoring tools that can do this; it’s that you want *all* your monitoring information *in one place*, whether it’s a simple Windows server or a completely-outsourced server or service.

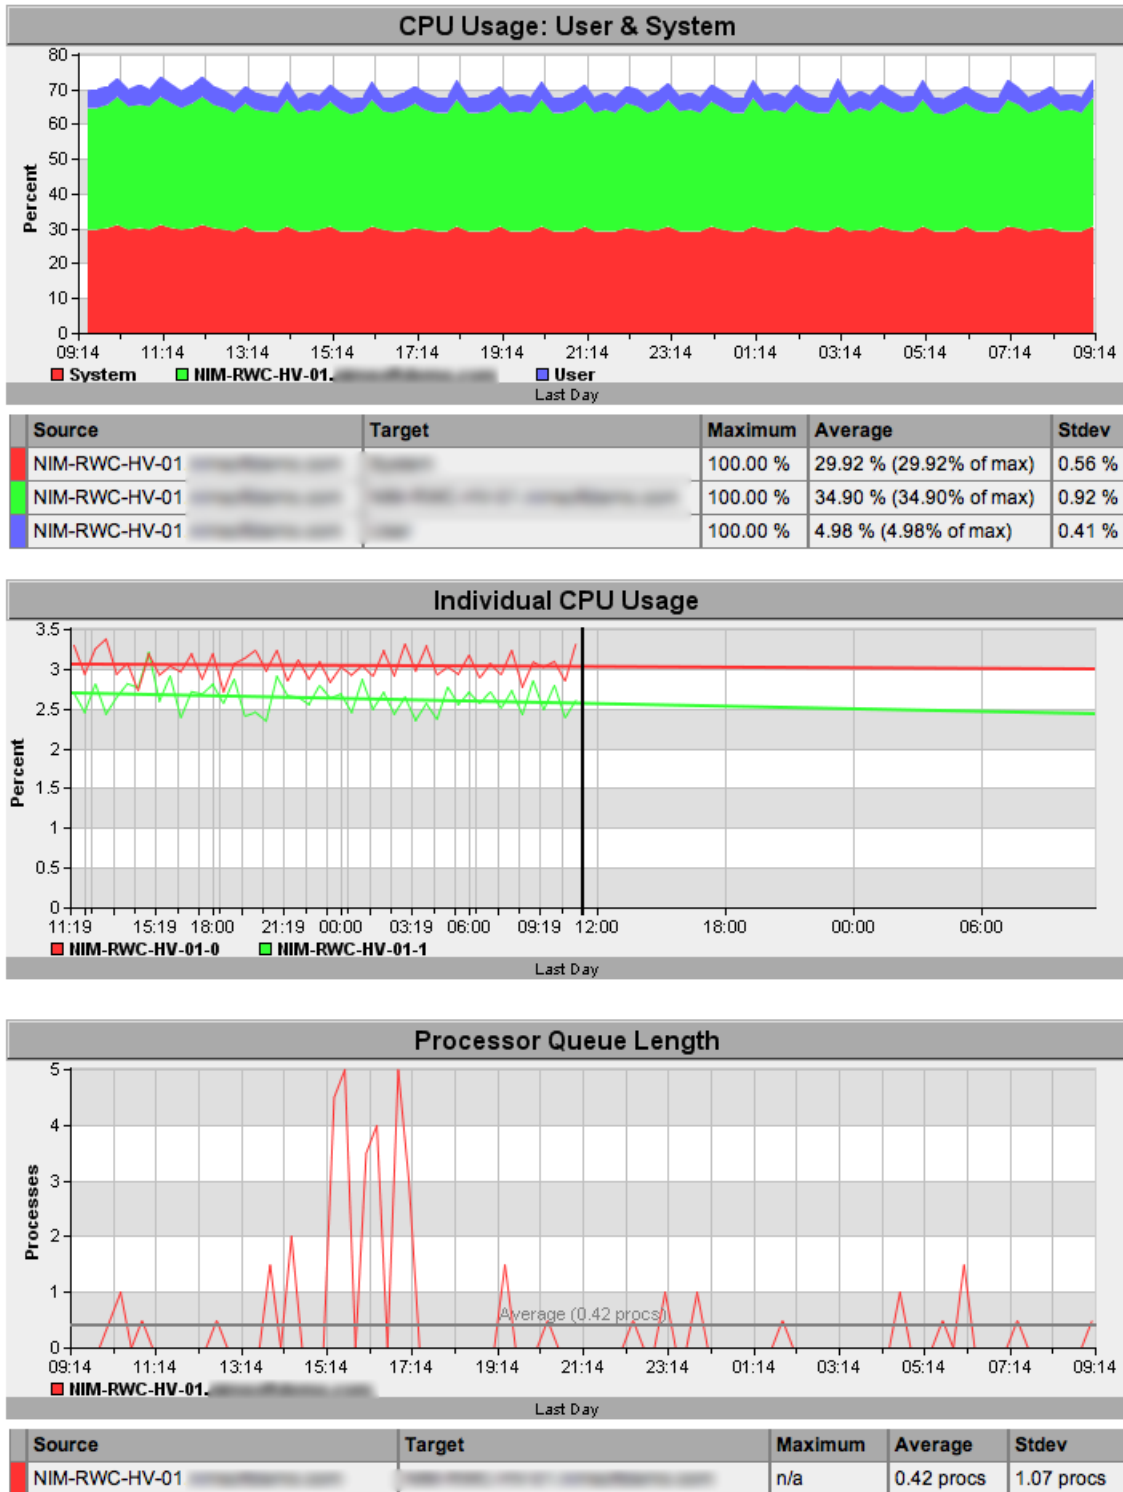
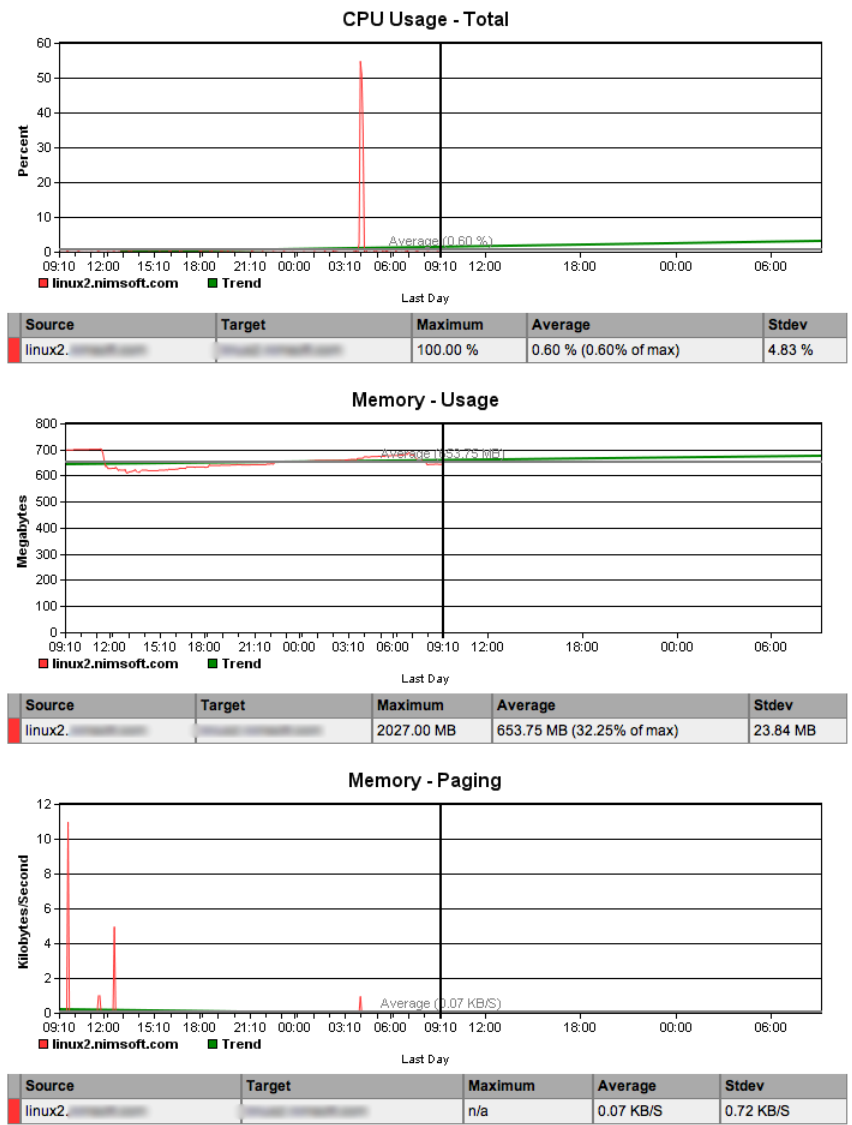


Figure 6.7: Windows server utilization.

**Note**

Figure 6.7 is huge—and I actually cropped off additional charts. This is one reason I prefer Web-based reports, because the browser can scroll as much as it needs in order to display large, detailed reports.

Broad platform coverage is a must. Even if you don't have Unix (or Linux) today, for example, you might well have a server or two in the future. As Figure 6.8 shows, your monitoring system needs to be able to accommodate that growth. I like to see reports like this, which essentially mirrors the Windows report and includes specifics for Unix. Windows and Unix are similar, and their performance would be monitored similarly, but a monitoring system can't ignore their unique aspects.



**Figure 6.8: Linux server utilization.**

As the use of virtualization grows, so must your ability to monitor it—no matter which brand you’re using. Figure 6.9 shows guest performance statistics on an IBM virtualization host—using terms and elements that are specific to IBM’s implementation.

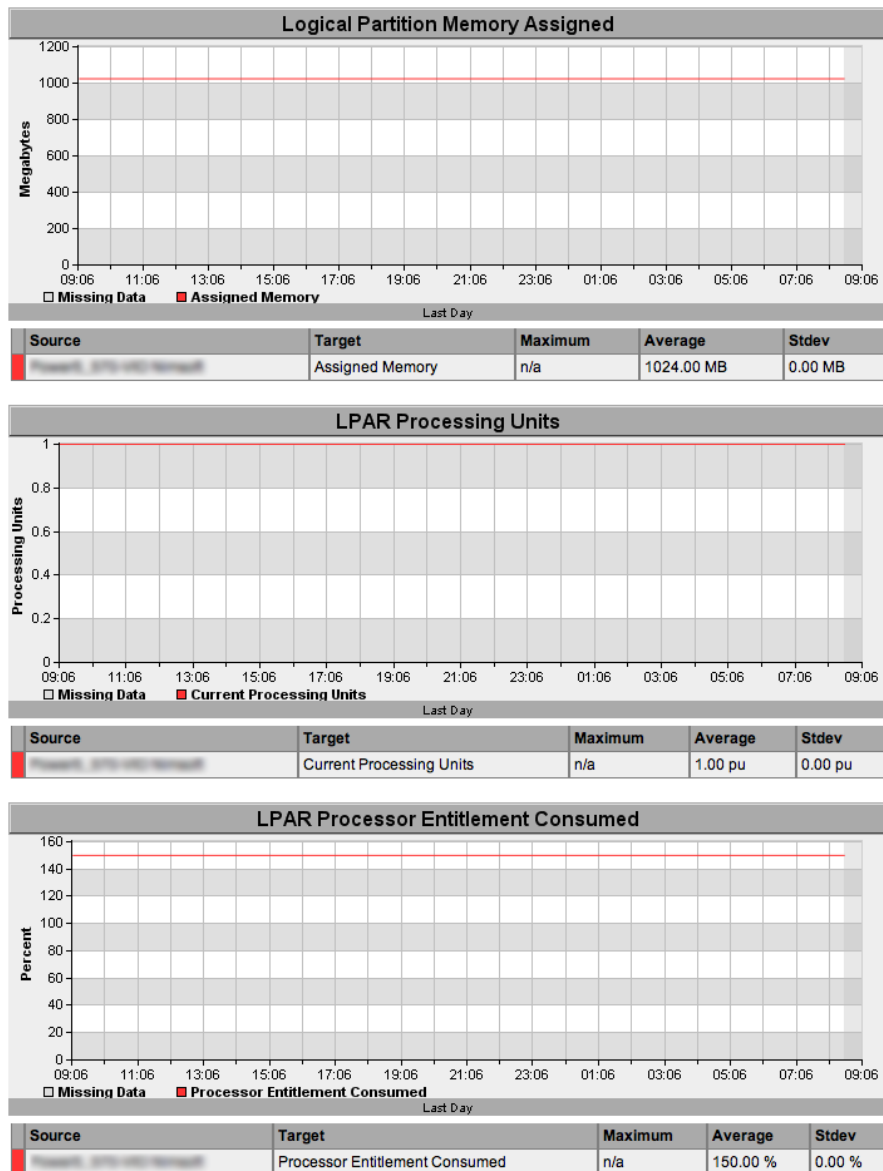


Figure 6.9: IBM virtualization guest statistics.

**Note**

Another excellent server—notice the stable performance over time.

Figure 6.10, however, shows that a monitoring solution can include other brands—such as VMware. This figure focuses on host statistics, showing key performance indicators for CPU, network, and so forth. Again, you want cross-platform reports to look similar so that you can do a sort of “apples to apples” mental comparison, but you don’t want to exclude vendor-specific information.

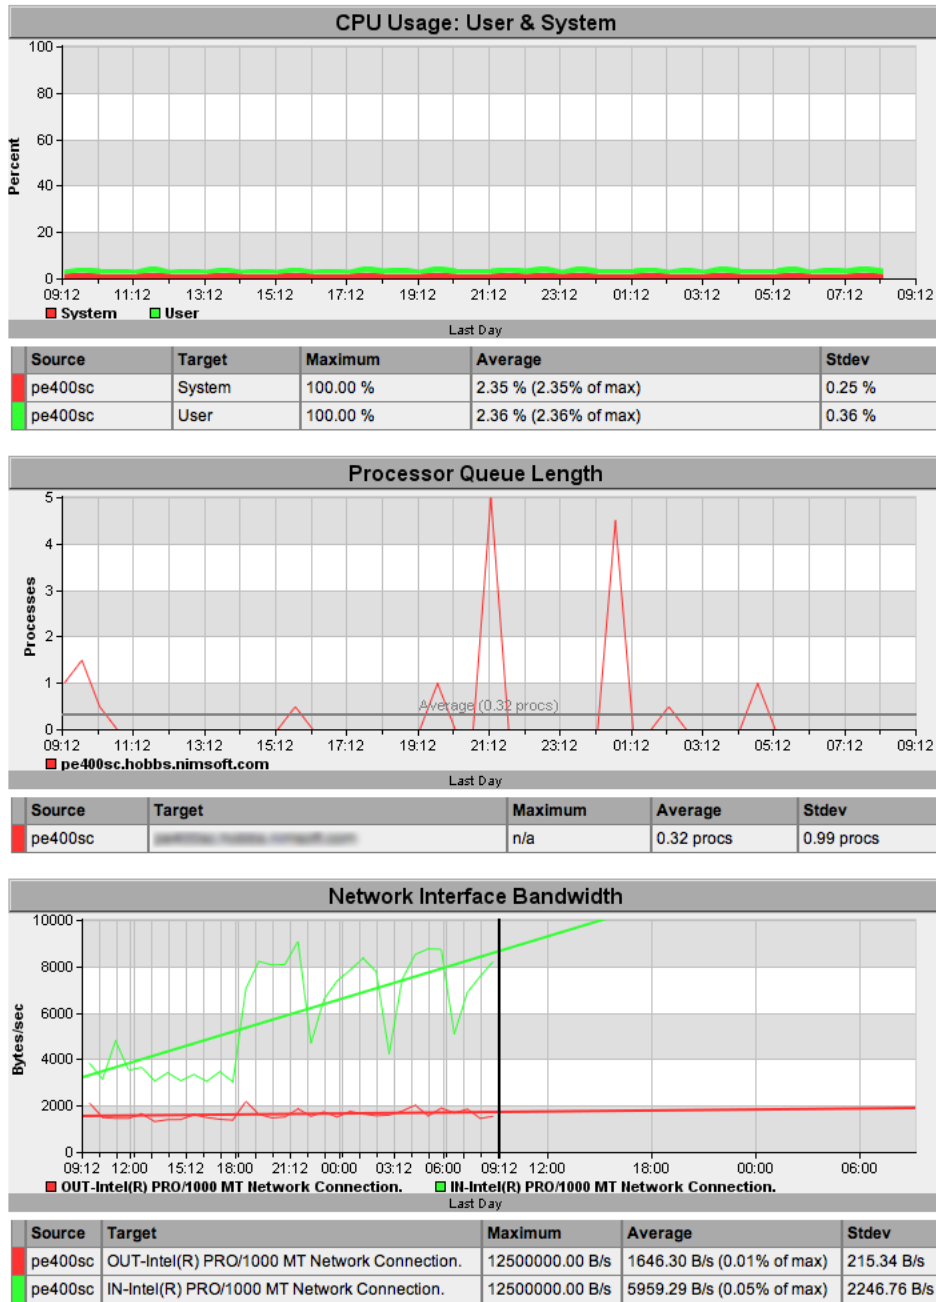


Figure 6.10: VMware vCenter monitoring.

More and more companies are adding VoIP to their technology mix, and there's no reason a monitoring system can't include that. Figure 6.11 shows a report for Cisco's CallManager system, providing a way to monitor and troubleshoot VoIP performance.

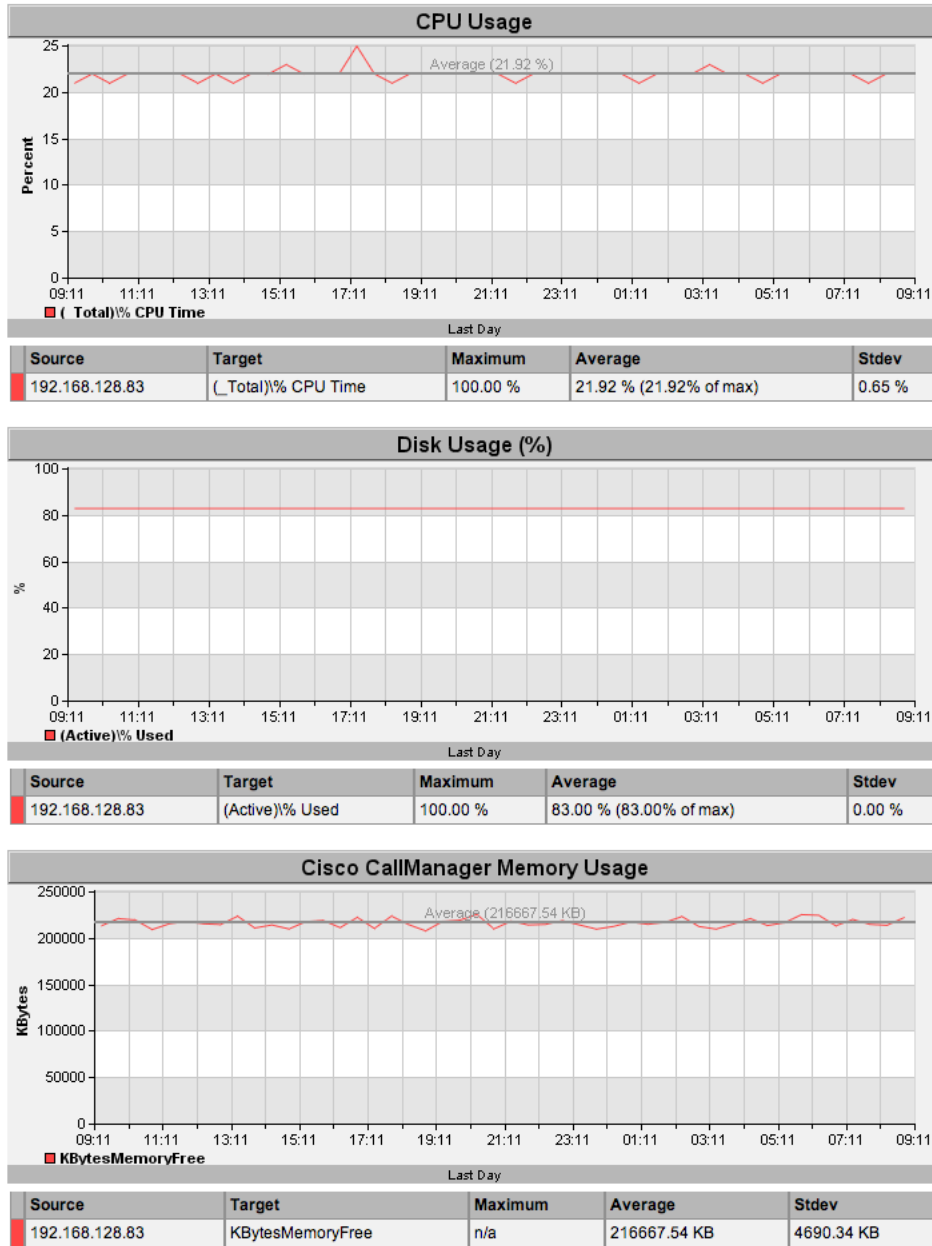


Figure 6.11: Cisco CallManager monitoring.



Every business has databases, and some of your applications will access those via Java Database Connectivity (JDBC), so you need to be able to monitor JDBC. Figure 6.12 is my first example of low-level, under-application monitoring, showing JDBC statistics. This might not be a report you look at first when a problem arises, but the point is that a monitoring system should provide this kind of information to help you dive deeper into a problem and either confirm or eliminate potential systems and technologies as the source of, or contributor to, a performance problem.



Figure 6.12: JDBC statistics.

You'll want to be able to monitor the database connectivity as well as the database platform itself. Figure 6.13 shows an example of MySQL monitoring, but your monitoring solution should include support for all major platforms, including Microsoft SQL Server, IBM DB2, Oracle, Sybase, and so on.

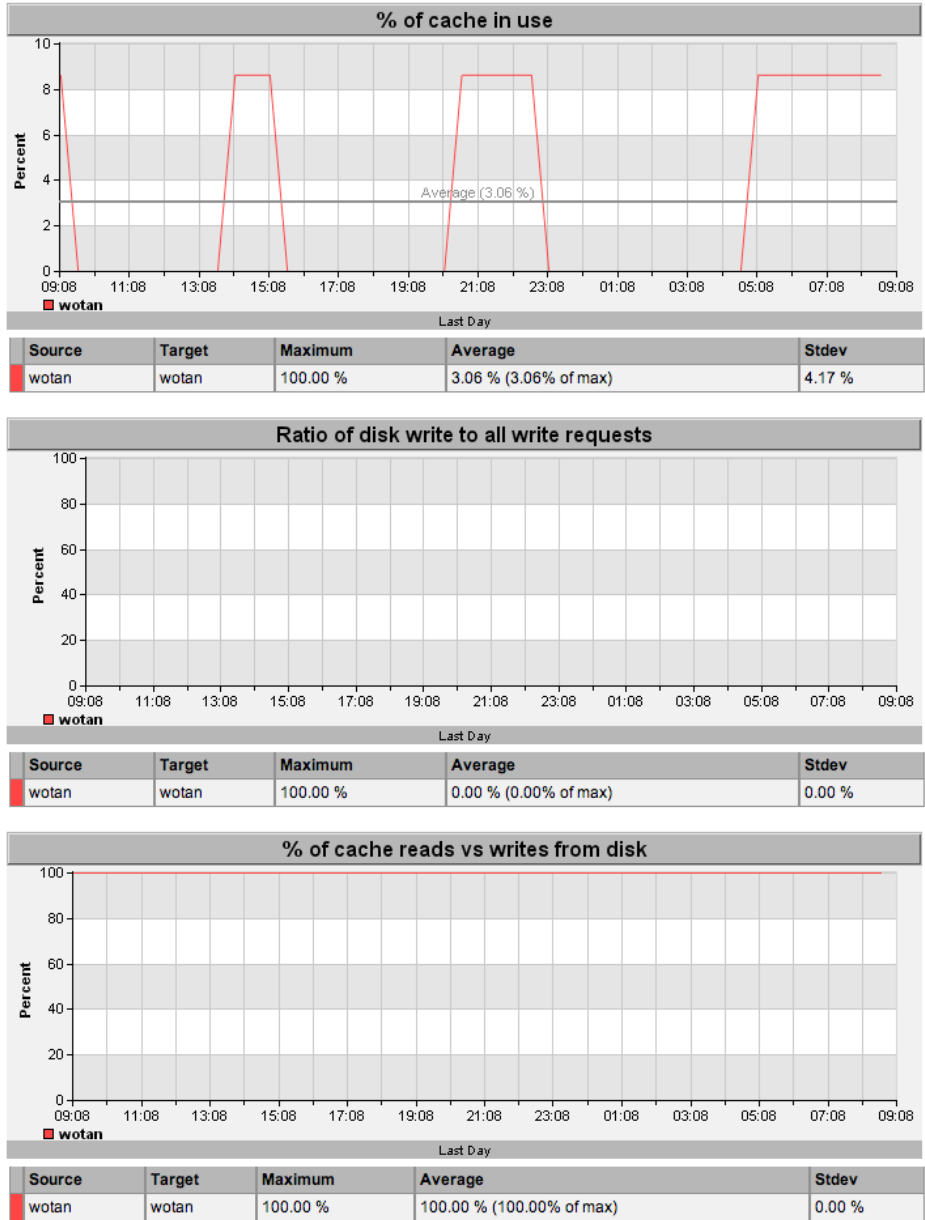


Figure 6.13: MySQL statistics.

Getting back to service levels for a moment, take a look at Figure 6.14, which shows how a monitoring system can also provide high-level, service-focused information—such as email round-trip times. This is a good indicator of overall system health, and I especially like the inclusion of a trend line that shows where performance is heading. That's a great way to get on top of a problem before it becomes a problem.




Figure 6.14: Email response times (SMTP).

Other services contribute to your overall IT performance, such as Active Directory (AD)—a lynchpin for many Microsoft-based (and third-party) services. Figure 6.15 shows that AD responsiveness can be monitored right within the same monitoring solution, watching statistics like connect time, replication speed, search load, and so on. Again, notice the inclusion of an average line, which lets you visually ignore peaks and valleys and focus on the overall average performance of a given service.




Figure 6.15: AD response times.

Web servers are running more and more applications, both internal and external, and monitoring the Web server platform is critical. Again, having this in a *single* solution makes it easier to monitor the entire application stack: Web server, database platform, database connectivity, network protocols, and so forth. You can start to see how this kind of system gives you insight into every aspect of the application, making it easier to spot and solve problems. Figure 6.16 looks at Microsoft’s IIS Web server.

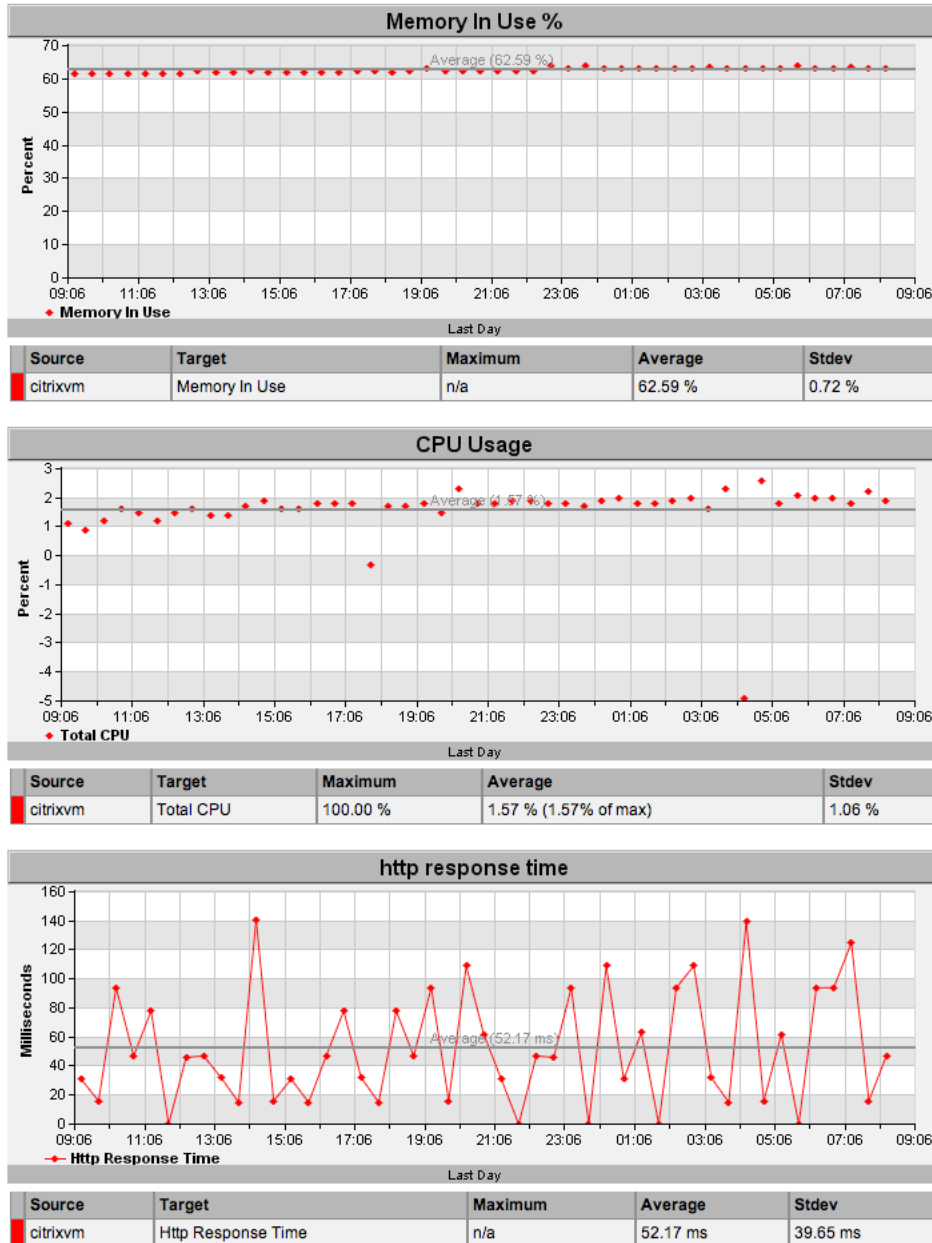


Figure 6.16: IIS Web Server statistics.

Because few shops are completely homogenous these days, Figure 6.17 shows that the same solution can also monitor the Apache Web server. Again, this report is similar to the IIS one, as both IIS and Apache are quite similar, but the Apache report includes specifics to that platform.

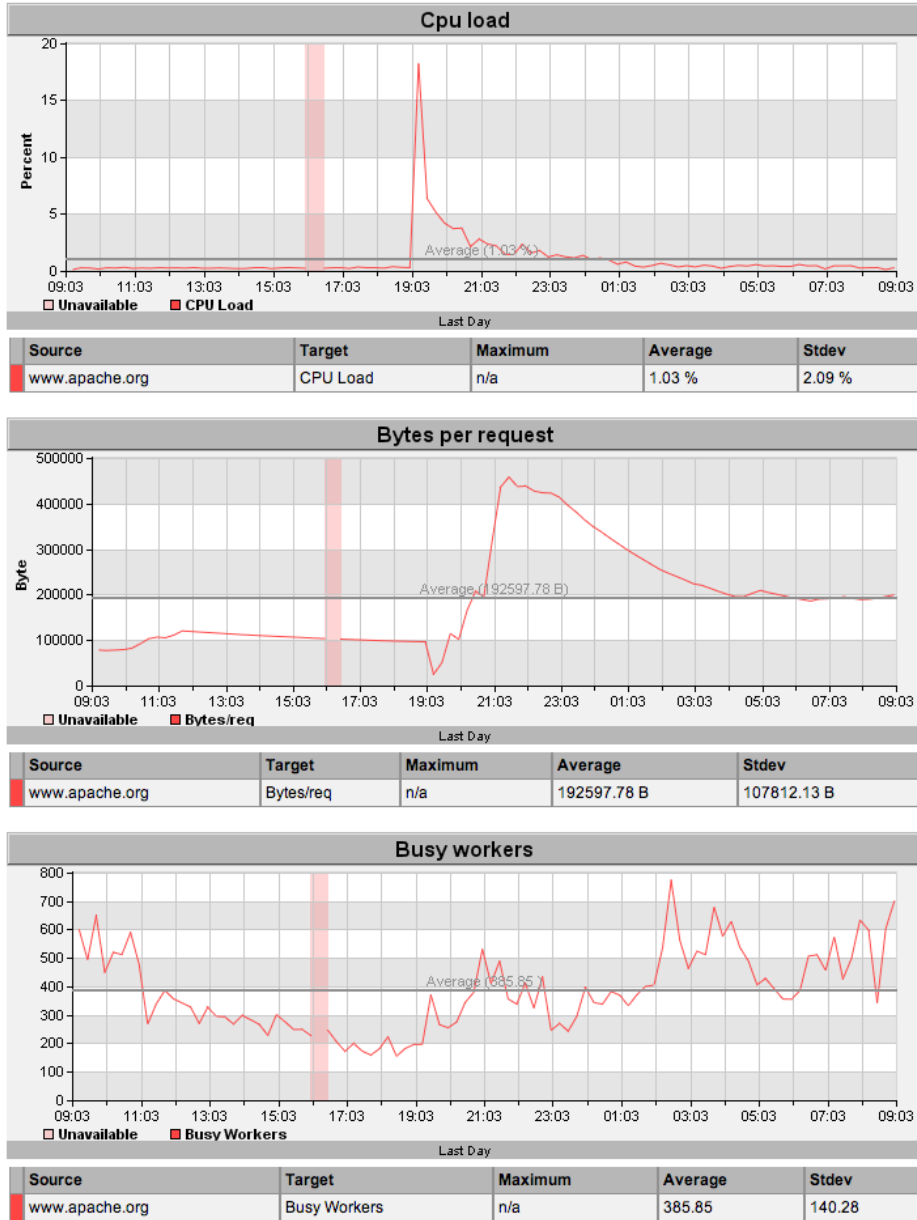


Figure 6.17: Apache Web Server statistics.

**Note**

See that vertical red area on each of the three charts? That's a time period during which my monitoring system wasn't able to talk to the Web server being monitored, so it couldn't draw the chart accurately.

Figure 6.18 once again returns to a service level-focused report, showing responsiveness for a Customer Relationship Management (CRM) solution. In fact, this particular report shows the results of synthetic transactions injected into the system to manage real-world performance from the end-user perspective—the *end-user experience* (EUE), that I’ve discussed in prior chapters. Here, we can see real-world response times for end-user tasks such as opening the application’s home page, logging in, and searching. A problem at this level would drive us to dive deeper—into the Web platform, database platform, network utilization, and so on.

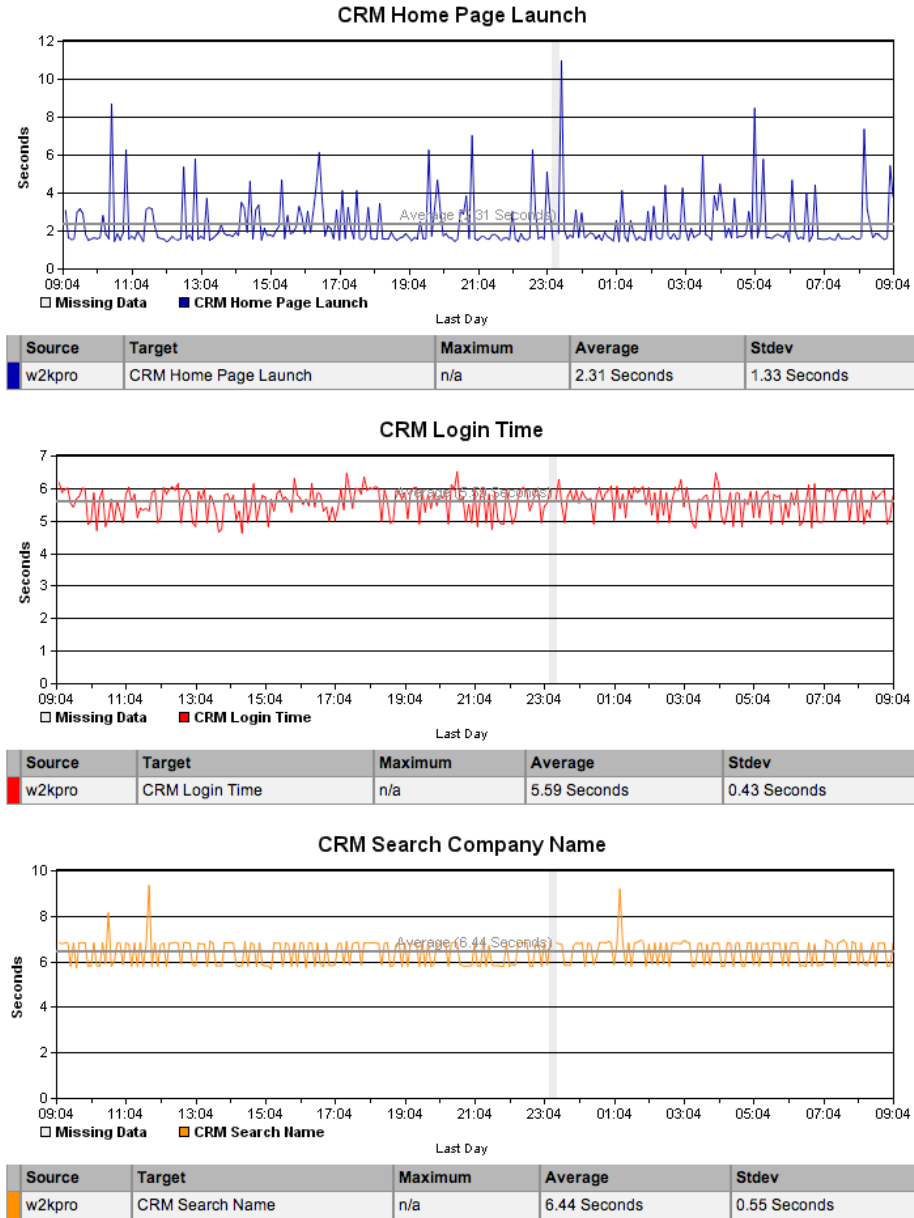


Figure 6.18: CRM system responsiveness.

Finally, you can't ignore the physical aspect of your infrastructure, and Figure 6.19 shows that a monitoring system can include considerations such as your server room's temperature—provided, of course, you have the right measurement probes in place to gather this information.

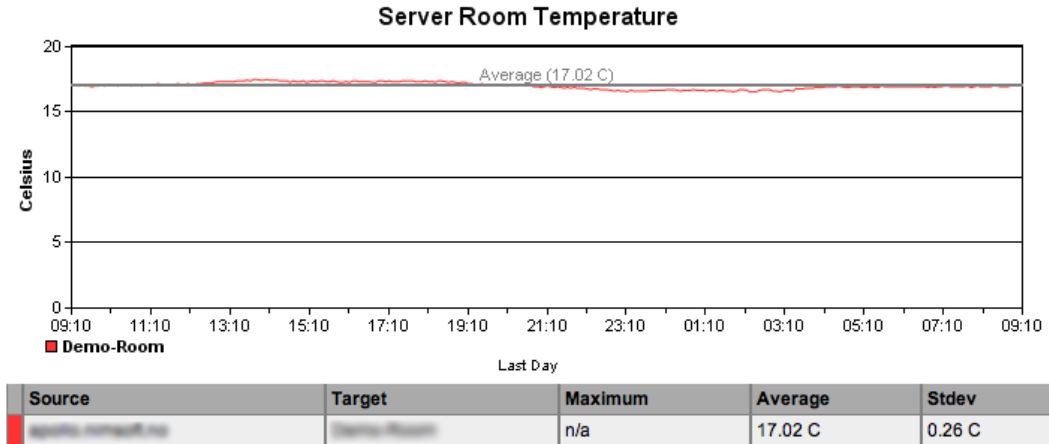


Figure 6.19: Server room temperature.

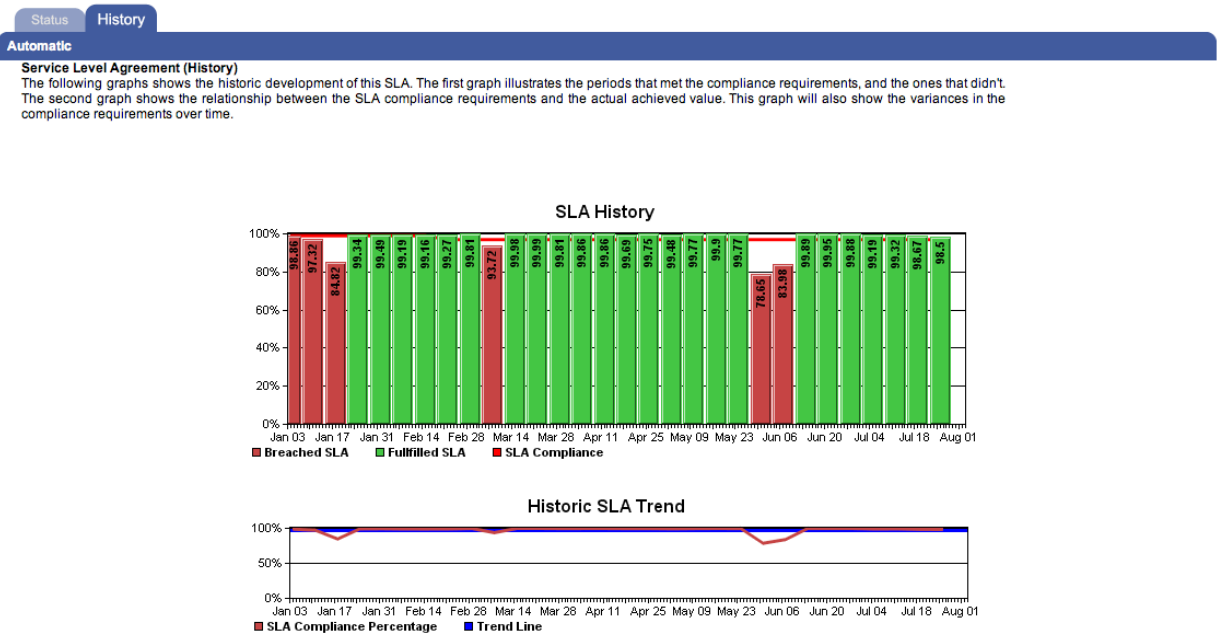
### SLA Reports

Having inundated you with examples, I'll just provide a couple for this section. The idea here is to roll up performance into something that can map to your SLAs, making it easier for you to manage those SLAs. Figure 6.20 shows the first example, rolling up numerous statistics into a simple "you made it or you didn't" measurement for several services, including a database server, Web server, Web response times as measured from two locations, and so forth. There's a trend analysis, too, indicating that the SLA is in no danger of being breached given current performance.



**Figure 6.20: An SLA report.**

A historical look is also nice, and Figure 6.21 shows an example.



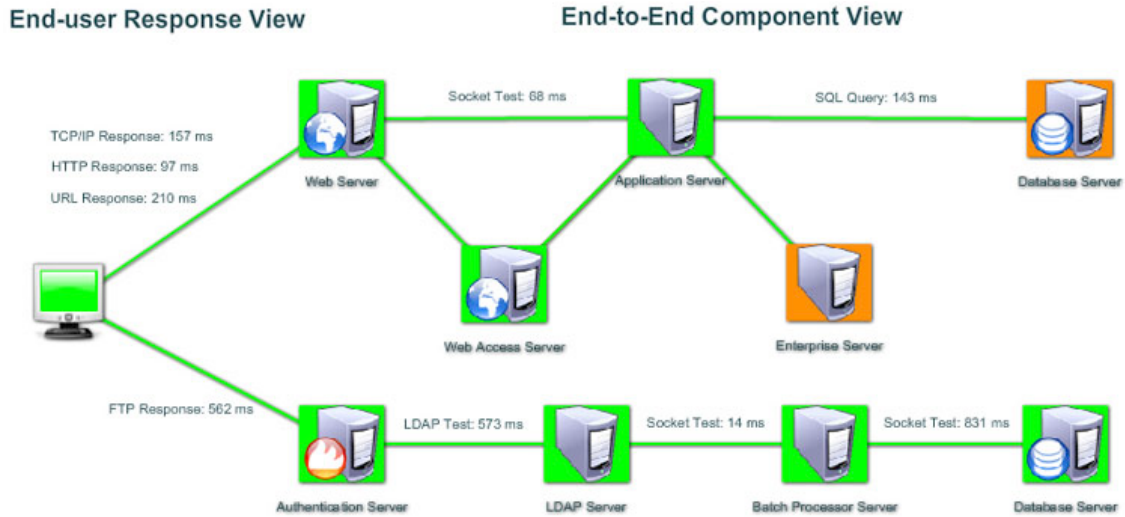
**Figure 6.21: Historical SLA performance.**

Here, we can easily see when the SLA was breached and how badly. This can be useful when it comes time to negotiate pricing or performance, especially for hosted services and applications.



## Dashboards

I love dashboards, and I dislike monitoring solutions that don't provide lots of 'em. These are a great tool for quickly checking the status of your environment at a high level, and for starting the detail dive when something is wrong. Figure 6.22 provides an excellent example, showing the end-to-end component view of an application, including individual servers, connectivity between them, response times for specific services such as SQL or LDAP queries, and so on. Problem systems are conveniently highlighted in orange, directing my attention to the components that require it.



**Figure 6.22: End-to-end performance dashboard.**

Figure 6.23 is an EUE dashboard, showing me—in simple colors and graphs—what my users are experiencing when they use a particular application (in this case, my Bugzilla bug-tracking application). I can see how fast the home page is launching, how long it takes to log in, how long it takes to find bugs and open them, and so on. I get a quick view (on the right) of the major platforms that comprise this application: the application code, MySQL database, Apache Web server, and Linux operating system (OS).

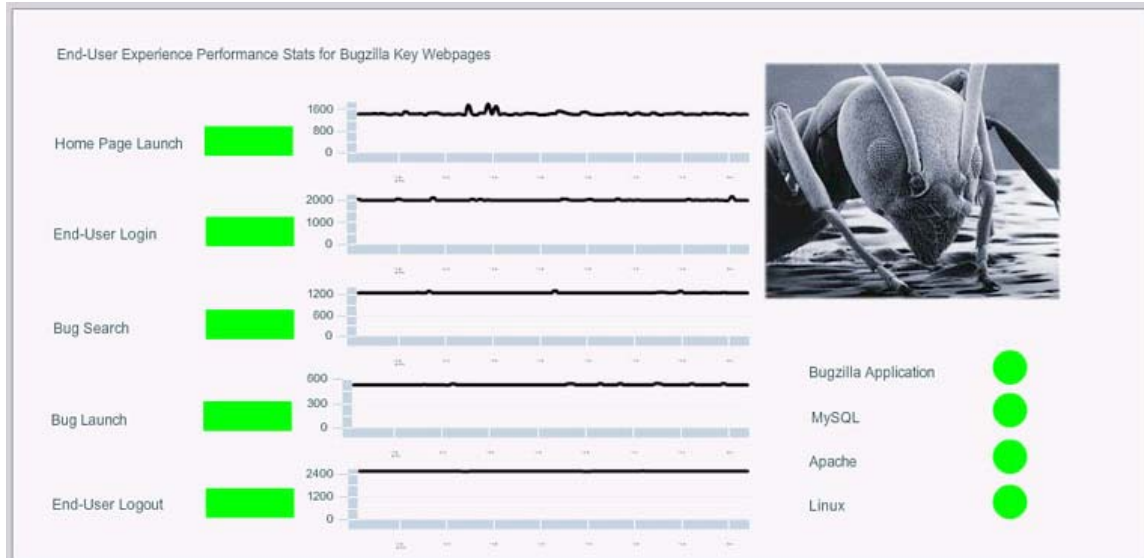


Figure 6.23: EUE dashboard.

For highly-distributed applications, geo-views like the one in Figure 6.24 are tremendously useful. This helps you see, at a large scale, where your application may be having specific problems.

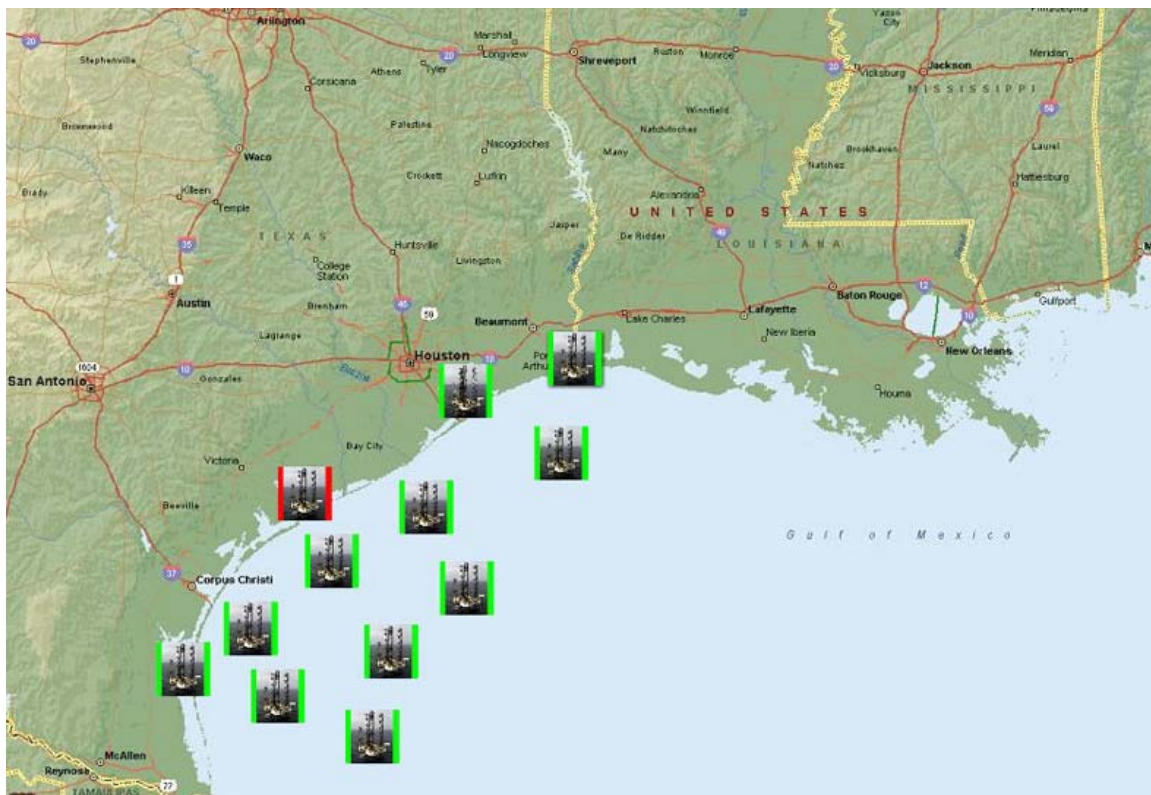


Figure 6.24: Geographic dashboard.

Finally, I especially like monitoring systems that can provide a customized, whole-environment rollup like the one shown in Figure 6.25, which was developed for a hospital. This dashboard provides an at-a-glance view of everything critical to healthcare applications in the environment. It’s literally the thing you want running all the time on some monitor somewhere so that everyone can be assured that all the systems are okay—or quickly take action if something isn’t.



Figure 6.25: High-level applications and services dashboard.

## The Provider Perspective: Reports for Your Customers

Managed Service Providers (MSPs) will appreciate most of the reports and dashboards I’ve shown so far, but they also need something specific to the kind of business they’re in. Most MSPs will also need the ability to look at performance from a client perspective so that they can see how a given client’s services are performing. A monitoring solution should absolutely be able to provide that, and Figure 6.26 shows one way in which it might do so: grouping services by customer and showing the overall utilization of each customer.

**MSP Master - ALL CLIENTS - Top Resource Consumers**

Contact Name: MSP Administrator  
 Number of Clients: 2  
 Servers Managed: 5  
 Network Devices Managed: 4  
 Total Devices Managed: 9

Customer	Hostname	Alarm Status	% CPU	% Memory	OS Platform	OS Version	OS Description
MSP	mzp-nimsrvr	Clear			Windows	Windows Server 2003 E	Service Pack 2 Build 3
ClientA	clienta-bugzilla	Clear			UNIX	Linux	Linux 2.6.18-6-686 #1
ClientA	clienta-nimsrvr	Clear			Windows	Windows XP Professic	Service Pack 3 Build 2
MSP	mzp-sum-vm	Clear			Windows	Windows XP Professic	Service Pack 3 Build 2
ClientB	clientb-nimsrvr	Clear			Windows	Windows XP Professic	Service Pack 3 Build 2
ClientB	clientb-xchsrvr	Minor			Windows	Windows XP Professic	Service Pack 3 Build 2
ClientA	clienta-websrvr	Minor			Windows	Windows XP Professic	Service Pack 3 Build 2

**Top 10 Servers by Resource**

Host	Interface	%In Value	%In Utilization	%Out Value	%Out Utilization
RWC_4500	Ethernet3/0/15		48.54		19.00
193.71.55.250	FastEthernet0/4		46.69		26.96
192.168.128.1	Ethernet3/0/15		39.24		19.92
RWC_TPAC_Router_1	ppp 1		15.18		4.94
RWC_TPAC_Router_1	t1 1/1: T1/FT1 Network Interface		15.11		4.92
RWC_TPAC_Router_1	t1 1/2: T1/FT1 Network Interface		15.10		4.91
192.168.128.1	Ethernet3/0/5		11.43		36.86

**Figure 6.26: MSP dashboard.**

## Conclusion

There you have it: Advanced, modern monitoring for the hybrid IT environment—from the data center to the cloud and even for MSPs that need to provide customers with insight into their own networks and systems. It *is* possible, using the right tools and the right techniques—and some vendors can even provide you with these monitoring capabilities as an SaaS solution, giving you an almost instant implementation, if desired. The solutions are out there—time to start looking.

## 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>.