



Best Practices for Maintaining a Healthy Exchange Messaging System

by Paul Robichaux

Contributing Editor, Windows & .NET Magazine

With any business, if a critical infrastructure system such as email degrades or fails, the effect will ripple outward to other business-critical operations and services. Microsoft has engineered Exchange to be both *robust* (not likely to fail, even under unusual conditions) and *resilient* (capable of graceful recovery when failures do occur). Despite these efforts, monitoring and managing your Exchange servers, and the infrastructure services they depend on, is critical to keep your overall business operations healthy. Your email system is more than just Exchange: it includes server and network hardware, the Windows OS, storage subsystems, and other components whose failure can affect messaging operations.

Of course, it's not enough just to blindly monitor *everything*: you have to intelligently monitor the things that are important. Smart monitoring will let you quickly assess the health of your messaging system, give you early warning when conditions change, and help you build a long-term plan to improve the service quality of your messaging and network services.

This paper will describe several things that you can, and should, monitor to ensure the health and service quality of your messaging system. The key to keeping your business operating is to proactively monitor key indices of system performance and health so that you can quickly respond to incipient problems. By watching these indices, you'll be able to better maintain the quality of service that your end users expect.

→ Contents

<u>Introduction</u>	1
<u>The Why and What of Monitoring</u>	1
<u>Ten Things to Watch</u>	1
<u>Interoperability</u>	2
<u>Operational Management</u>	3
<u>Security</u>	6
<u>Conclusion</u>	6
<u>About the Author</u>	6

Best Practices for Maintaining a Healthy Exchange Messaging System

Introduction

It's been said that the business world is a jungle. As it turns out, most companies' internal business operations are a jungle, too: an interdependent set of ostensibly separate pieces, some competing and some cooperating. When any component of the jungle ecosystem has trouble doing its job, the overall health of the ecosystem suffers. So it is with business: if a critical infrastructure system such as email degrades or fails, the effect will ripple outward to other business-critical operations and services.

Microsoft has engineered Exchange to be both *robust* (not likely to fail, even under unusual conditions) and *resilient* (capable of graceful recovery when failures do occur). Despite these efforts, monitoring and managing your Exchange servers, and the infrastructure services they depend on, is critical to keep your overall business operations healthy. Your email system is more than just Exchange: it includes server and network hardware, the Windows OS, storage subsystems, and other components whose failure can affect messaging operations.

Of course, it's not enough just to blindly monitor *everything*: you have to intelligently monitor the things that are important. Smart monitoring will let you quickly assess the health of your messaging system, give you early warning when conditions change, and help you build a long-term plan to improve the service quality of your messaging and network services.

This paper will describe several things that you can, and should, monitor to ensure the health and service quality of your messaging system. The key to keeping your business operating is to proactively monitor key indices of system performance and health so that you can quickly respond to incipient problems. By watching these indices, you'll be able to better maintain the quality of service that your end users expect. (See "Ten Things to Watch" for a list of the 10 most critical items to monitor.)

The Why and What of Monitoring

Have you ever called a business and heard a recorded announcement saying, "for service quality, your call may

be recorded"? This recording notifies you that the business is monitoring the quality of service that it provides—and the same principle is useful for Exchange systems. The goal of setting up monitoring is to allow you to set and track meaningful performance and service quality goals. The users who depend on your messaging systems to get their jobs done have certain expectations, and so do the other line-of-business systems (and users!) who depend on the messaging system. Most organizations have set up service-level agreements (SLAs) that define exactly what levels of service are required—but someone has to monitor the service to make sure that everyone involved is getting what the SLA promises.

The exact type and metrics of SLAs for different environments will vary; they might include the average time required to deliver a message to an internal or external recipient, the average time required to open a new message, or the

Ten Things To Watch

If you're in a hurry to get started, you may be wondering what are the most critical items to monitor. Here's a brief list; each of these items is covered in more detail later in the paper.

Interoperability Monitoring

1. Message flow
2. DNS availability and response time
3. Active Directory availability and replication health

Operations Monitoring

4. Exchange storage performance and availability
5. Exchange and Active Directory CPU usage
6. Network availability and performance
7. Key event occurrences (-1018 errors, other critical indicators)

Security Monitoring

8. Network traffic analysis and monitoring
9. Security event monitoring
10. Intrusion detection

maximum amount of downtime allowed per year or per month. Some SLAs focus solely on end-user metrics (“Newly submitted messages must be delivered to internal recipients within 5 minutes”), while others are oriented toward more traditional measures (“All inbound messages must be virus scanned; this scanning should not delay inbound messages more than 5 minutes.”)

How do you know what to monitor? In general, you want to measure whatever values, parameters, or actions will give you a way to tell how well your messaging environment is meeting your service requirements. (Actually *defining* those requirements is outside the scope of this paper.) Everything you monitor should have a connection to some aspect of your service quality or service level metrics. This connection might be indirect. For example, if your Exchange server runs out of disk space on the transaction log volume, the Information Store service will stop gracefully, rendering your server unable to send or receive mail or to accept client connections. This in turn affects the availability of applications and business processes that depend on email flow—so disk space utilization should be monitored. If you can’t make the connection between a monitorable parameter and your service quality, you probably shouldn’t be monitoring it.

Whenever possible, you’ll get a better overall “big picture” view of your messaging system if your monitoring tools allow you to collect multiple related statistical measures and combine them into a composite measurement. For example, if you’re interested in the amount of time it takes a message to reach the Internet, being able to easily see the queue length and average delivery time for one server is useful—but being able to see the average of these values for *all* servers at the same time (and then drill down to inspect individual servers when needed) is invaluable (Figure 1).

Exchange, Active Directory (AD), and Windows all expose a variety of measurements that you can monitor and analyze. Choosing the right things to look at is critical, so the remainder of this paper is devoted toward highlighting key items that you should regularly be watching as part of your Exchange deployment. These items fall into three general categories:

- **Interoperability**, which covers message exchange with the outside world and replication and message exchange within the enterprise-messaging network. Monitoring items in this category will tell you when mail or replication traffic slows or stops, whether internally or externally.
- **Monitoring and management**, which covers the operational monitoring of servers, the work they do, and the messages they transfer. Of course, this category also includes monitoring the performance and status of the network, and watching for critical events whose occurrence

means that something’s wrong.

- **Security**, which includes intrusion detection, auditing, and network traffic monitoring.

Interoperability

Exchange is often used as the backbone system for other business processes. For example, a financial research company might have processes built around generating and sending out research reports; a pharmaceutical company might use Exchange with a third-party document management system to manage and control the flow of regulatory and government approval-related documents. Even when Exchange is used only for “regular” messaging, the flow of those messages within the company, and with external recipients, is often critical to the company’s business.

The first, and obvious, group of things to monitor in this category involves message flow with the outside world. When your users aren’t getting mail, is it because no one’s sending mail, or because there’s a problem with mail flow? To find answers to such questions, you’ll want to monitor:

- The total number of inbound and outbound messages sent per some time unit (per-minute and per-hour are both useful scales); this allows you to see when the flow increases or decreases suddenly.
- The health and status of virus scanners, archiving tools, or other add-on programs that process incoming or outgoing mail; for example, if you see that your virus scanner is using an unusually large, or small, amount of CPU time over a sustained period, that’s a possible red flag.
- The number of failed or rejected inbound or outbound messages. Spikes in these numbers may indicate spam attacks or communication problems with other sites.
- Queue behavior, including the average queue length for key connectors between routing groups or between your organization and the outside world. This helps you understand how long it takes a new message to be processed for delivery, and that’s a good indication of how busy your email servers and support systems are.

Of course, email-related measurements aren’t the only interoperability-related parameters you need to watch: parameters that indicate the state of AD and DNS services are critical because failures or slowdowns in these services will affect your messaging, and in turn your business processes. You should consider monitoring:

- **Replication flow**. Because Exchange depends on AD, the point of monitoring inter-site replication between AD domain controllers is to ensure that changes made to the

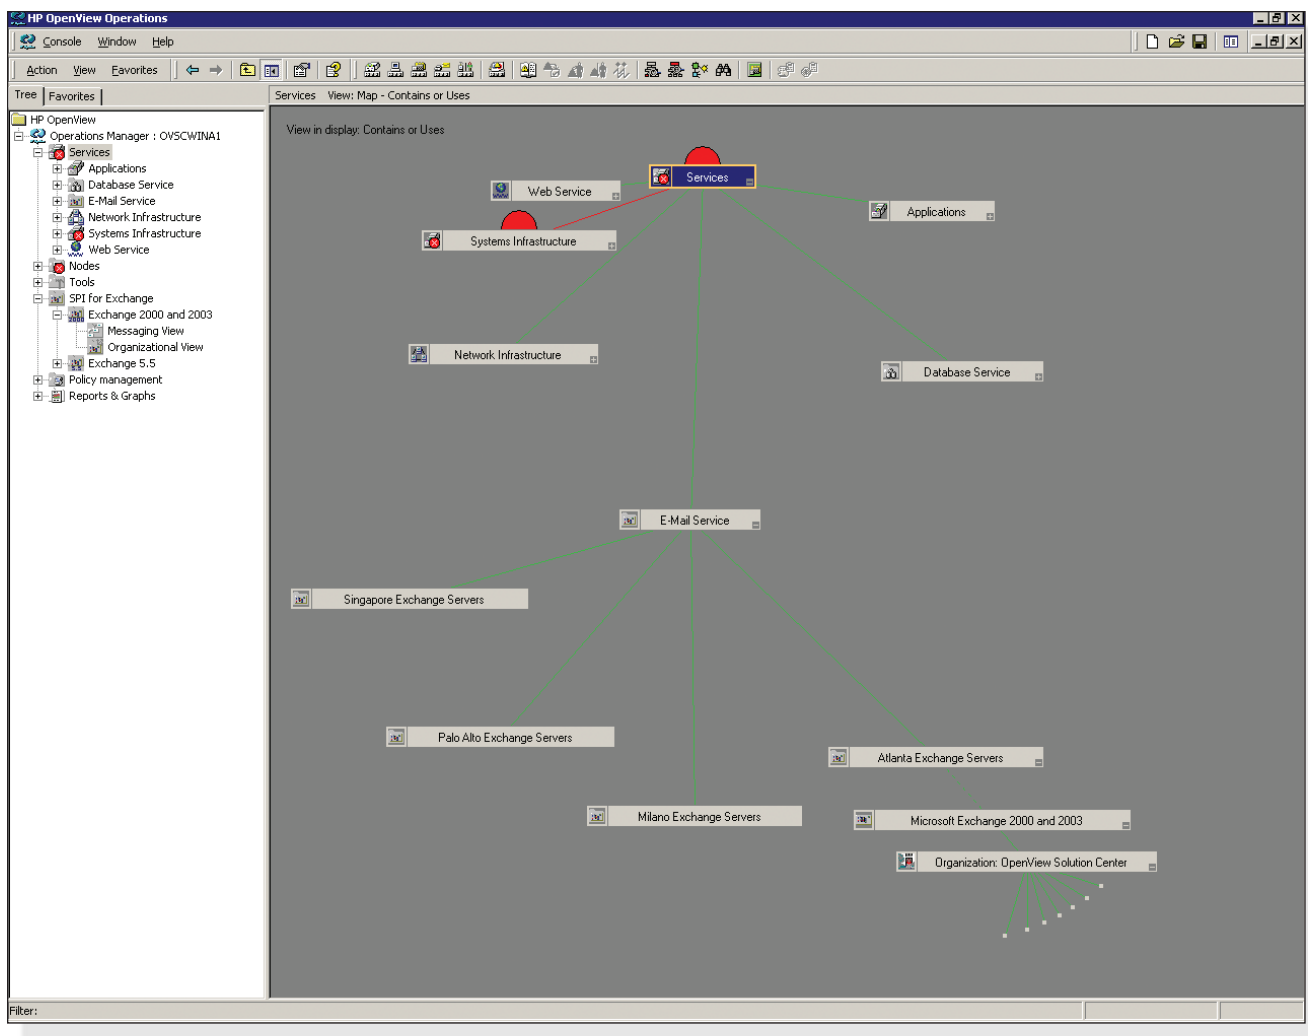


FIGURE 1
Getting the “big picture” for Exchange

directory—like adding and removing user accounts, creating or removing groups, or changing Exchange configuration settings—are properly replicated. Careful monitoring of the latency and response time of replication requests (Figure 2) will tell you whether your domain controllers and global catalogs are keeping up; monitoring the replication latency time can also alert you to a class of problems nicknamed “creeping death,” a situation in which replication gradually slows to a halt because of inconsistencies in replicated data or errors caused by network or server problems.

Operational Management

Once you’ve set up monitoring to tell you how well your messaging systems are interoperating with each other and the out-

side world, you’re ready to set up monitoring for operational management. This category actually encompasses much of what Exchange administrators think of as their daily tasks; however, the things you monitor here actually go beyond Exchange into the underlying health and performance of Windows and your network’s infrastructure services.

The first thing to consider monitoring is storage utilization and performance. Most Exchange administrators are familiar with the groundbreaking performance analysis work done by Pierre Bijaoui’s team at the HP Technology Leadership Group, and the recommendations that have come from it. For example, you should monitor the average disk queue length for the volumes that hold your Exchange databases; the average queue length shouldn’t exceed the number of physical disks in the volume. For a single disk, that means an average queue length

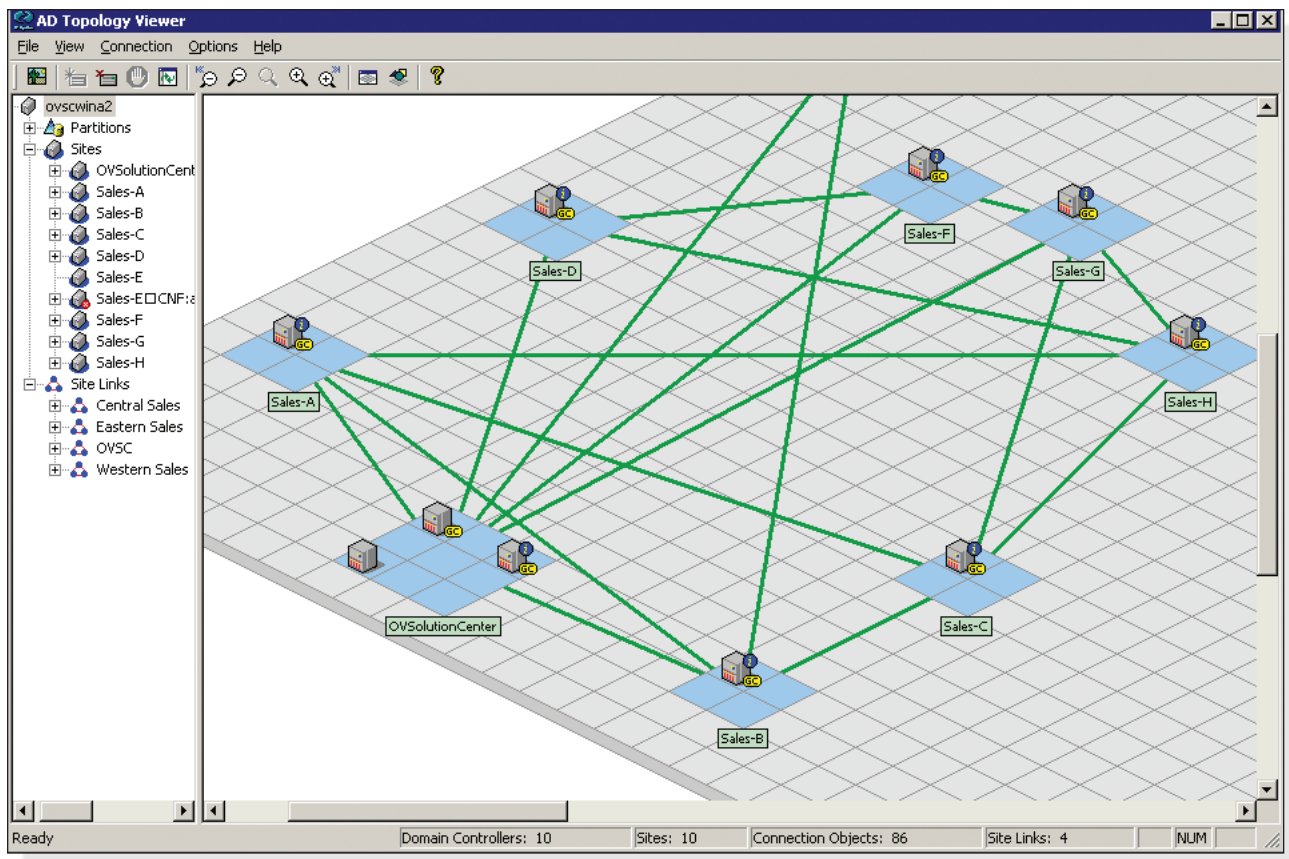


FIGURE 2
Monitoring AD replication and topology health

less than 1; for a five-disk RAID-5 array, that means a queue length less than 5. Of course, monitoring storage *utilization* is critical too, because running out of space is a good way to shut down your Exchange server. You'll want to be able to quickly see how much space your Exchange servers are using and how much free space remains on each server; in addition, it's helpful to have a tool that can alert you when the rate of storage consumption suddenly goes up—that alert may enable you to stop a message loop or other problem before it affects your operations.

New Exchange storage technologies demand their own kind of monitoring, too. For example, if you're using a hardware or software replication solution to provide redundancy for disaster recovery, you should monitor the replication transport (whether it's SAN- or network-based) to ensure that your data are replicated with acceptable latency. If you're using a network-attached storage (NAS) device with Exchange, that typically means you'll have a private interconnect between the Exchange server and the NAS server—

and that interconnect needs to be monitored, too.

Of course, storage isn't the only thing that needs to be monitored. To do a good job of operational monitoring, you have to include measurements of the other resources Exchange needs to work properly, including the performance of Exchange and AD services. You can measure these resources in two ways: by tracking the amount of CPU time consumed by those services, or by tracking the amount of time it takes particular tasks (such as on-demand virus scans or brick-level backups) to complete. This gives you a good idea for how well your servers' CPUs are bearing up under everyday load, and unexpected changes in CPU usage may signal a deeper problem.

Network performance and service quality monitoring are very important to monitoring the quality of your email service. That's because if the network is down, or slow, everything from replication to DNS queries to sending mail between servers in a routing group to getting up-to-date antivirus signatures may be negatively affected. The first thing that users notice, of course, is when mail doesn't get where it's supposed to in a

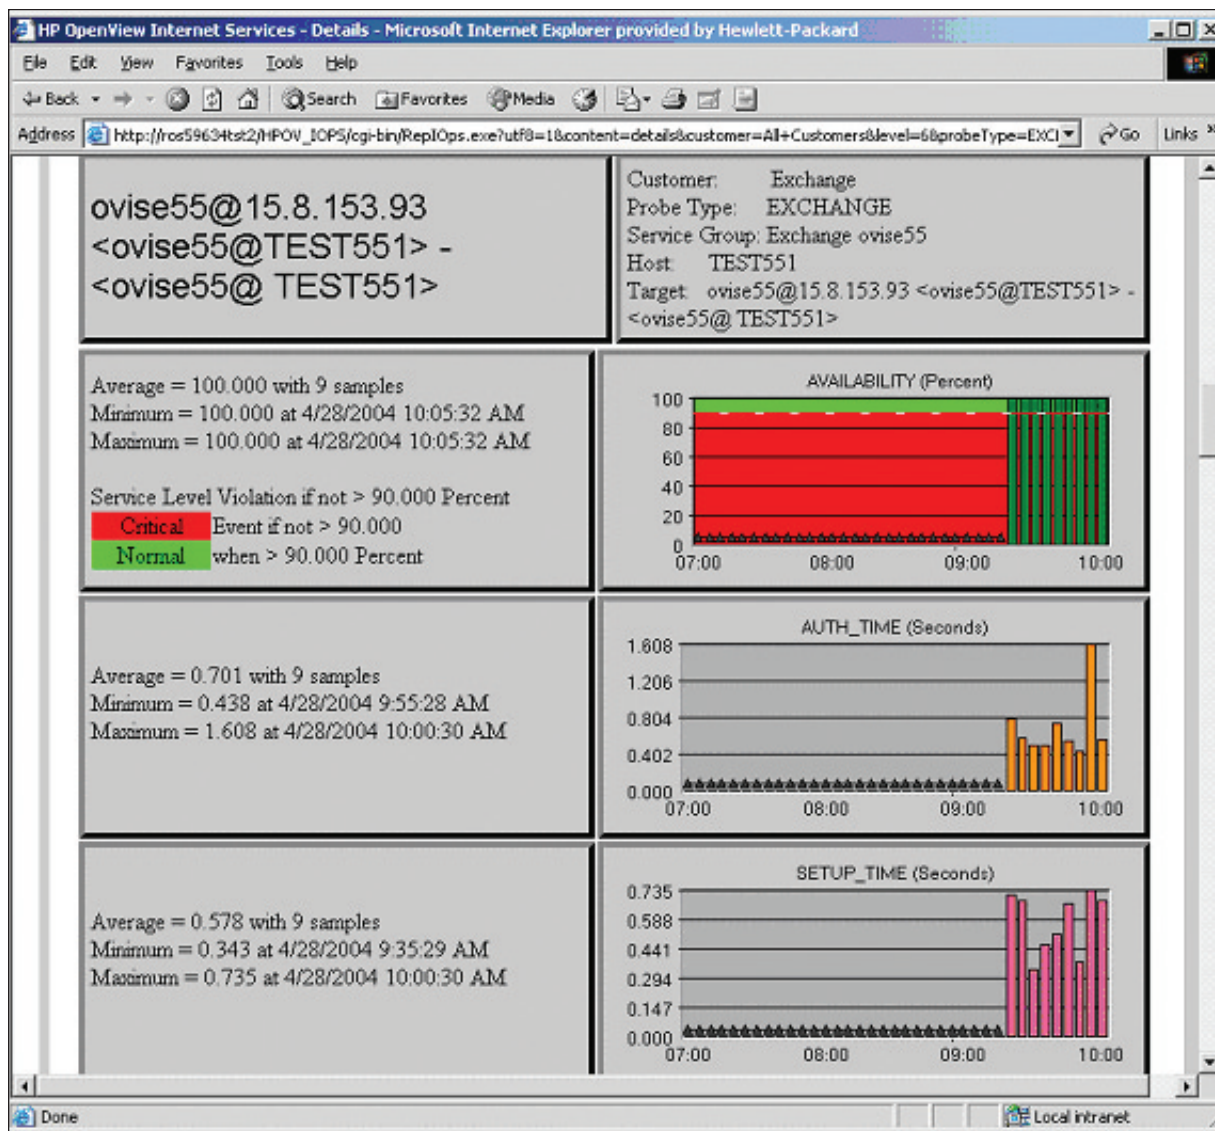


FIGURE 3
Monitoring the Exchange user experience

reasonable time. So watching the network for signs of congestion or outage can help you keep on top of problems.

As an example, management tools such as HP's OpenView can synthesize aggregate measurements from several lower-level measurements to give you a good picture of what the end-user experience looks like. For example, watching the amount of time it takes to accept and deliver a new message—something that users expect to happen quickly—tells you when things are running slower than normal. Monitoring the response time required for DNS queries, or the time required to get a directory update from one AD site to another, can

alert you to low-level infrastructure problems that can affect Exchange, and other business processes that depend on it, if left unfixed (Figure 3).

The last operational monitoring concern is event monitoring. Some events generated by Exchange and Windows require immediate attention. For example, the infamous MSeXchangeIS-1018 event, which indicates physical corruption to a database page, is something you'll always want immediate notification about. Likewise, your backup program probably posts event log messages when a backup error occurs, and you'd want to find out about those as soon as possible. As you consider monitoring

solutions, be sure that you have a way to specify conditions for which you want immediate alerts.

Security

Security monitoring builds on the other types of monitoring discussed earlier. In many large organizations, a dedicated security team is responsible for ensuring physical and electronic security for company assets. However, at most organizations, network and messaging administrators have to secure the network themselves, and this begins with being aware of the state of the network. Things you'll want to be able to monitor include:

- The appearance of certain kinds of servers on the network. For example, it's a good idea to monitor the presence of DHCP servers, because a rogue DHCP server can be used as a component of a broader network-borne attack. If you see new DHCP servers popping up on the network, this might be a sign that something's wrong.
- The presence, absence, or traffic flow for certain kinds of protocols. With the right real-time networking monitoring tools, you can easily be alerted when users use protocols you don't want on your network (like BitTorrent, Gnutella, or Napster), or when unusual traffic begins to rise in volume. One case where this can be very useful is with viruses and worms that cause a compromised workstation to act as an SMTP server: if you suddenly see an upsurge in SMTP traffic from machines that aren't Exchange servers, you should find out why. Most worms have a characteristic network signature—machines infected with the Slammer worm, for example, could be identified by watching for ICMP packets with a specific, and unusual, length. A system that lets you quickly add rules to detect certain types of network traffic can be very valuable as an additional security layer.
- The appearance of certain kinds of events. For example, Windows can be configured to generate success and failure

events for logon attempts, password changes, and other security-related occurrences. Watching for patterns—like a string of several logon failures, followed by a logon success—can be very useful for ferreting out improper activity.

- Intrusion detection. Most intrusion detection systems (IDS) are complex, purpose-built tools that can monitor a wide range of security-related parameters and behaviors. However, an integrated management tool such as OpenView can tie the IDS into the rest of your measurement and monitoring system to provide you with a clear picture of what the IDS is finding.

One caution, though: remember that monitoring for security may require you to monitor things that don't have any connection to your quality of service measurements. Don't let these data distract you: focus on the aggregate indications that tell you what's *really* going on.

Conclusion

Many administrators realize that Exchange is important to their everyday business operations, but sometimes it takes a prolonged outage to reveal exactly *how* important. Because Exchange is often the backbone of other business processes or services, an unexpected service interruption may have more serious consequences than just keeping employees from mailing each other. However, these consequences can be kept at bay by careful monitoring. This monitoring leads to proactive problem management, and that in turn translates directly to lower operating costs, better user satisfaction, and less work for administrators. ◆

*Paul Robichaux is a principal engineer for 3sharp and an MCSE. He is the author of several books, including *Secure Messaging with Exchange 2000* (Microsoft Press), and creator of the <http://www.exchangefaq.org> Web site. Paul writes regularly for *Windows & .NET Magazine*.*