

THE TOP 10 MISCONCEPTIONS ABOUT PERFORMANCE AND AVAILABILITY MONITORING

ABSTRACT

Performance and availability monitoring are essential for ensuring the health of the enterprise's mission-critical applications. But as technology has quickly evolved, some of yesterday's monitoring best practices are no longer as valid. Today's complex, distributed applications are leading companies to reevaluate their beliefs about monitoring and to add new monitoring best practices to their strategies.

This white paper will examine 10 common misunderstandings about monitoring and offer suggestions for implementing complete, end-user monitoring strategies that can adapt to today's more complex computing environments.

TABLE OF CONTENTS

Abstract	2	6. Monitoring all of the Available Metrics for a System or Application is the Best Approach	5
Introduction	3	7. Second or Sub-second Sampling Rates are Necessary	5
The Top 10 Misconceptions about Monitoring	3	8. The Best Monitoring Solutions Include Built-in, Automatic Corrective Actions	6
1. Monitoring Basic Infrastructure is Enough	3	9. The Monitoring Software has to Reside In-house	6
2. The Only Way to Monitor is with an Agent	3	10. A Company's Infrastructure Monitoring Strategy Can Operate in its Own, Detached Silo	6
3. One End-user Experience Monitoring Technology is Better than Another	5	Mercury Business Availability Center	7
4. Monitoring Processes or Services for an Application Suffices	5	Summary and for More Information	7
5. All Systems in the IT Enterprise Must be Monitored	5		

Introduction

Today's enterprises depend on the availability and performance of their mission-critical business applications. If these applications suffer from degradations in performance or fail completely, companies are subject to lost revenue and decreased customer satisfaction. In order to avoid these undesirable outcomes, IT departments must adopt effective monitoring strategies without actually making problems worse or breaking the bank in terms of total cost of ownership (TCO).

Since the earliest days of business computing – with monolithic mainframe installations and very small client/server-based systems – IT departments have well understood the importance of monitoring availability and performance and have invested in various forms of monitoring solutions. But as technology has rapidly evolved, some of yesterday's best practices are no longer valid. In some cases, following outdated strategies can lead to ineffective monitoring, high overhead, and increased costs. In addition, today's complex, distributed applications on which many mission-critical business processes rely are forcing companies to reevaluate their long-held beliefs about monitoring and to add new best practices to their existing strategies.

The purpose of this white paper is to draw attention to 10 aspects of monitoring that may have once been perfectly sound best practices, but now, due to new technologies and/or changes in IT infrastructures, may no longer apply in most availability and performance monitoring situations.

The Top 10 Misconceptions about Monitoring

1. Monitoring basic infrastructure is enough.

Monitoring system metrics (such as CPU, memory, and disk) is important – but these metrics do not provide enough information to truly understand whether actual users or applications are experiencing performance problems. Trying to “add up” individual system performance metrics in order to understand actual application or end-user performance doesn't work either.

True end-user-oriented monitoring is critical and should be the starting point of any monitoring strategy. In addition, due to advances in hardware reliability and performance, the causes of most performance problems today are usually problems with application components, as opposed to individual pieces of hardware. Thus, system monitoring alone – while still critical – will not provide an accurate or complete picture of true application performance.

2. The only way to monitor is with an agent.

Infrastructure availability and performance monitoring initially started with deploying heavyweight agents that added significant overhead to the production systems, and therefore introduced the possibility that the agents could either crash the system or become the cause of the very problems they were trying to help avoid. Companies lived with this possibility because few other options existed.

Today there is another way – agentless monitoring. Agentless monitoring enables operations groups to monitor complex, distributed systems without installing agents or software on the production systems. These solutions non-intrusively monitor any and all parts of the IT system remotely from a single host machine. Agentless monitoring solutions ensure rapid deployment; ease application maintenance and upgrades; reduce the risks of impacting production systems; facilitate system infrastructure and expansion; and increase return on investment (ROI) by decreasing the TCO of the monitoring infrastructure.

Some companies have hesitated to implement agentless monitoring due to several common misconceptions. The most common agentless misconceptions are:

– **Agentless monitoring requires too much bandwidth.**

Agentless solutions connect remotely to the monitored production servers and do create some network traffic. However, in order to monitor most basic system and application metrics for the purposes of identifying sustained performance problems, very little bandwidth is needed. For example, monitoring basic CPU, memory, and disk information takes little more than 10KB total per sampling interval per system. With today's corporate networks regularly deploying 100MB bandwidth, more than enough capacity exists.

– **Agentless monitoring isn't secure.**

Agentless monitoring can be as secure as IT administrators want it to be. Since most agentless monitoring solutions simply login to a remote system just like any user would, login security can be maintained by using appropriate security steps and credentials for the user defined for monitoring. In many cases, access can be granted only for the system components or resources that need to be monitored.

In addition, many agentless solutions have the ability to further secure remote monitoring via technologies such as Secure Shell (SSH) and HTTPS. For example, Mercury's agentless monitoring solution, Mercury SiteScope®, can be configured to establish remote connections via SSH for Windows, UNIX, and Linux remote systems, and can also transmit data back to the Mercury SiteScope server via HTTPS.

– **Agentless monitoring doesn't provide enough granularity.**

Another misconception is that agentless monitoring cannot collect enough data or granular enough data to effectively monitor mission-critical systems. While it is true that certain low-level metrics may not be easily accessed by agentless monitoring, metrics that provide the ability to identify the most important monitoring issue – sustained performance problems – are easily gathered.

It should also be noted that in the past few years, many technology vendors have embedded interfaces (such as JMX, XML, SOAP, etc.) into their products that enable a wide variety of data to be gathered using agentless monitoring. Finally, research indicates that collecting too much data can dramatically drive up the costs of monitoring and in most cases, data collected from metrics other than "key indicator" metrics is rarely used.

Many enterprises have already purchased an agent-based infrastructure monitoring framework from one of the "big four" management vendors (Hewlett-Packard, IBM Tivoli, Computer Associates, and BMC). Although these solutions may be providing some value to the enterprise, most rely on high-cost, agent-based technologies. Therefore, whether it is the high cost of maintenance, high overall TCO, or simply extremely long implementation times, enterprises that rely solely on these solutions are not monitoring everything they'd like to monitor.

In addition, the "big four" agent-based solutions tend to be much more infrastructure-, system-, or network-focused at a time when application monitoring and end-user experience analysis is critical. Finally, solid integrations between the "big four" and other lower TCO monitoring solutions can enable an effective co-existence strategy that makes sure the entire enterprise is covered.

3. One end-user experience monitoring technology is better than another.

When it comes to monitoring the end-user perspective, much debate has raged about the accuracy and overhead of different end-user monitoring approaches. The fact is that several end-user monitoring technologies are required, depending on the situation. There are three categories of end-user monitoring technologies – business process monitors, real-user monitors, and client monitors.

- **Business process monitors** run synthetic transactions to capture the performance and availability as experienced by end users. Synthetic monitors use agents distributed throughout the Internet to *simulate* how a particular website or application will react to heavy loads and provide information about application performance from distributed geographies. Synthetic monitoring solutions enable enterprises to drive active transactions against applications from outside the company firewall to monitor business processes externally.
- **Client monitors** sit on the actual end user's machine and capture and report on the performance and availability of real transactions executed by the end user.
- **Real-user monitors** sit off of the network node to capture and report on the performance and availability of URLs and transactions. These monitors measure application performance for many existing end users by tracking actual user traffic. Real-user measurement tools often use appliances or software probes to passively monitor all client interactions by attaching to a mirror port on the edge router in a data center.

All three technologies are important and each one has a role to play when creating a complete end-user monitoring solution.

4. Monitoring processes or services for an application suffices.

Today's applications – whether they be packaged applications, J2EE-, or .Net-based – are complex and span multiple systems and various technologies. Simply monitoring a few key services or processes will not provide a complete picture of application health and certainly will not provide the level of detail needed to troubleshoot thorny performance problems. In order to thoroughly understand application health, component monitoring is required.

5. All systems in the IT enterprise must be monitored.

While it is tempting to monitor everything that uses electricity in the IT enterprise, 100-percent coverage is not necessary. IT enterprises typically consist of several systems that don't support business-critical functions or applications. The trick is in knowing which systems relate to critical business functions and which ones don't. Application relationship mapping technology can help ensure that IT is monitoring the systems, applications, and application components that really matter.

6. Monitoring all of the available metrics for a system or application is the best approach.

Performance problems tend to follow Pareto's Rule – "80 percent of problems are generally caused by 20 percent of the system's or application's components." The challenge is in knowing which metrics are the "key indicators." Otherwise, either too much data is collected or the wrong metrics are monitored. Instead of monitoring every possible metric, IT administrators should look for monitoring solutions with built-in expertise regarding the most important metrics to watch.

7. Second or sub-second sampling rates are necessary.

The most important alerts needed when monitoring infrastructure performance and availability are the ones for sustained performance problems. Monitoring with second or sub-second intervals is not

necessary to identify sustained performance issues and usually results in massive amounts of data that is never used, or “alert storms” that trigger too many people getting involved in a situation that may not be an emergency.

Second or sub-second monitoring is also likely to uncover events that are temporary or transitional, and not necessarily good indicators of performance problems that really impact end-user experience. While it is true that some aspects of performance and availability may execute faster than a second and therefore require sub-second sampling, these are few and far between, especially when it comes to basic infrastructure.

8. The best monitoring solutions include built-in, automatic corrective actions.

While corrective actions can be useful in a very few application performance situations, rarely do automated corrective actions see the light of day in most companies' monitoring strategies. Few system administrators are willing to trust a software tool to take action on its own.

Secondly, hardly any corrective actions (besides server reboot, temporary file cleanup, etc.) can be applied as the remedy for the types of performance problems that are showing up in current distributed applications. Most performance and availability problems are a result of application component issues. A custom, scripted approach is generally the best strategy for taking corrective actions, especially if it provides the ability to control whether it runs automatically or not.

9. The monitoring software has to reside in-house.

Application management outsourcing has gone mainstream. The reason for this is generally due to its lower TCO, faster implementation time, and ability to deliver insight without end users or even administrators needing to become application monitoring gurus. In addition, an outsourced monitoring strategy offers the ability to “test” a company’s infrastructure from outside of its firewalls, preferably from multiple locations around the world.

10. A company’s infrastructure monitoring strategy can operate in its own, detached silo.

Today’s enterprise monitoring strategies are becoming more and more tied to other strategies within the organization. Pre-production testing and development need feedback from the real-time monitoring teams when designing new applications or tweaking existing ones. Change management strategies and solutions must be factored in when determining the cause of performance issues. Business people must be involved in helping to set thresholds and service-level agreements (SLAs). The bottom line is that monitoring must be a good corporate citizen and integrate with a myriad of other solutions, strategies, and processes.

Today’s enterprises need to have the ability to ensure that their applications and systems are meeting their established performance and availability requirements in both pre-production and production. IT organizations must therefore have the ability to monitor, diagnose, and resolve critical problems across the entire application lifecycle. Effective performance optimization and management solutions must be able to span the entire performance lifecycle from development to production. Specifically, these solutions should enable IT to:

- Performance-test applications prior to rolling them out to production – mitigating the risks of application downtime.
- Use capacity planning to create the best architecture in the production environment – by optimizing across cost, performance, and utilization requirements.

- Monitor, measure, and manage enterprise applications and the underlying infrastructure in production.
- Proactively diagnose and resolve application problems in both test and production environments.

Mercury Business Availability Center

Mercury Business Availability Center™ is the only end-to-end solution that enables proactive business service management. It provides complete visibility into and control over the end-user status and business availability of applications running in complex, distributed application environments. Mercury Business Availability Center enables enterprises to map, measure, and manage application, system, and infrastructure performance and availability according to end-user, service levels, and business goals.

Mercury Business Availability Center helps enterprises quantify the business impact of application downtime and resolve performance problems when they arise. It also offers the ability to obtain real-time visibility into the complex and changing relationships between applications and the underlying infrastructure.

Mercury Business Availability Center includes integrated applications and a business dashboard for performance and application monitoring, agentless system availability management, service-level management, configuration management, application mapping, diagnostics, and problem resolution. As a result, IT departments can reduce mean time to identification (MTTI), improve service-level performance, minimize application downtime, and lower TCO.

Mercury Managed Services™ for Business Availability Center is a remote, outsourced monitoring service that enables enterprises to optimize the performance and availability of all internal and external-facing applications by leveraging Mercury's pre-deployed infrastructure, operations, and expertise. Companies can start with Mercury Managed Services for Business Availability Center and transition to an in-house implementation if they choose to at a later time. Mercury Managed Services offers the benefit of testing from multiple points of presence around the world and eliminates the need to employ monitoring gurus to implement, since it is all managed for the customer with minimal involvement.

Summary and for More Information

Application monitoring is essential for improving the uptime and availability of the enterprise's mission-critical, distributed systems. Effective monitoring solutions should provide the ability to protect the company's revenue streams and ensure that all systems are performing up to the company's pre-defined service levels.

But monitoring best practices have changed substantially from the earliest days of business computing. And as this technology has evolved, older monitoring practices may no longer be effective. Today's enterprises need to re-evaluate their existing beliefs about monitoring and adopt more innovative monitoring best practices to support their new application environments and strategies.

Mercury Business Availability Center provides a proactive, scalable, end-to-end approach to monitoring and managing business applications and systems. Mercury Business Availability Center ensures that mission-critical business processes are not only up and running, but performing according to the enterprise's most critical business requirements. For more information on application monitoring or any Mercury products and

MERCURY[™]

Mercury is the global leader in business technology optimization (BTO). We are committed to helping customers optimize the business value of IT.
WWW.MERCURY.COM

© 2005 Mercury Interactive Corporation. Patents pending. All rights reserved. Mercury Interactive, the Mercury logo, Mercury Business Availability Center, Mercury Managed Services, and Mercury SiteScope are trademarks or registered trademarks of Mercury Interactive Corporation in the United States and/or other foreign countries. All other company, brand, and product names are marks of their respective holders. WP-1491-0705