

Intelligent Monitoring of Distributed Applications:
11 Tips for Delivering the Performance
Your End Users Demand



Intelligent Monitoring of Distributed Applications: 11 Tips for Delivering the Performance Your End Users Demand

Abstract

Today's application end users—whether within your organization or outside it—are more demanding than ever. That's true whether the end user is a bank customer accessing an ATM, one of your sales managers checking last year's revenue, or a supplier logging onto your extranet to view the latest RFPs. They expect that the distributed applications they need will always be available and operating at peak performance.

It's the responsibility of the IT department, and specifically the IT manager, to deliver on those expectations. But, just as these expectations are rising—and systems infrastructures are becoming more complex—many of the resources available to IT managers, in the form of budgets, systems resources, and staff size, are static or shrinking. The challenge for these managers is to leverage their finite resources effectively to fulfill their performance commitments.

This white paper examines the challenges facing IT managers and the key role of application monitoring and management in meeting end-user expectations. It presents 11 specific tips for monitoring and managing your distributed applications to help deliver the performance levels demanded today and in the future.

Serving Today's Empowered End User

Welcome to the world of the empowered end user. As an IT manager today, you have a bigger burden in terms of deliverables than your predecessors did. It's no longer enough to keep your applications and infrastructure up and running. You still have to do that and manage your department for maximum productivity and economy.

But there's a new set of challenges out there. Now, more than ever, the empowered end user is a part of the IT service equation. First, the end users you are serving are no longer just employees within your organization. They can be your company's business partners and suppliers who interact via corporate extranet for a wide range of e-Business functions. Or they can be one of your company's clients, with whom you have formal service level commitments.

The ultimate end user may even be the person-on-the-street: the customer ordering an airline ticket from your e-Commerce site, the person calling customer service to resolve a purchase problem, or the supermarket shopper waiting to get through the checkout line.

What do these end users expect? High on the list are:

- 100 percent application availability, wherever and whenever they want it;
- Reliable communications and short response time across the network;
- Robust, high-performance applications that can handle their processing needs without slowing down or breaking down;
- Fast, easy access to information resources, whether they're in local, remote, or even external databases, or Web-based repositories;
- Full interoperability and transparency across new as well as legacy applications, operating systems, and networks; and
- The latest applications and features that will help them do their jobs—or accomplish any online task—better, faster, and more profitably.

But Wait, There's More: Meeting Your Own Management Goals

It would almost be easy if the end user's expectations were the only things you, as an IT manager, had to fulfill. But, as department head, you have a whole additional set of performance goals. You need to:

- Stay on budget, both for individual projects and your total IT operations, by keeping IT costs to a minimum;
- Make sure your staff is maximally productive;
- Leverage your IT staff's expertise to the maximum, so that, for instance, your most highly skilled systems managers and application developers don't focus on routine maintenance tasks; and
- Plan for the future, developing new IT strategies so that you don't fall behind the technology curve. That means researching, assessing, and implementing new technologies, integrating new and legacy applications, and continually adapting your infrastructure to meet your organization's evolving business needs.

The question for IT managers is how do you balance the need to meet end-user expectations with your own management goals? Either one could take 100+ percent of your time, attention, funds, and resources. The key to keeping end users' demands under control is to provide the necessary performance levels through solid techniques and best practices for monitoring and managing your distributed applications.

Evaluating Your Monitoring and Management Options

Given the complexity of computing environments, it's more challenging than ever to achieve and sustain peak performance from your distributed applications. The good news is that, thanks to the ongoing evolution of monitoring and management technology over the years, IT managers now have more—and more sophisticated—options than ever before.

By now, most IT departments are aware of the need for application monitoring and management and have taken one of three different routes in addressing the issue. They have either created their own proprietary solution, purchased and customized a packaged solution, or chosen to work with built-in tools or freeware.

How do you, as IT manager, determine which is the best for your organization and infrastructure? And, once you have a monitoring and management solution in place, how can you make sure you're getting the full benefit of that technology?

To answer these crucial questions, we first need to define some of the basic concepts in application monitoring and management. To monitor an application, you need a component that collects application performance data and evaluates it to identify real or potential problems. The component is usually called an agent, and the instructions that identify problems are usually called sensors or rules. The IT staff can view performance data from a central system, or console.

The exact “flavor” of a monitoring and management solution is a function of its structure. For instance, it may deploy local agents on each monitored system or rely on SNMP to collect performance data. Further, it can employ intelligent local agents that not only collect and evaluate data but can take independent action when events occur or pre-defined performance thresholds are met. That action can be either automated (e.g., if an application fails, the agent automatically initiates a restart) or non-automated (e.g., the agent alerts IT staff to the problem and they decide on the appropriate action).

Other ways in which monitoring and management approaches can differ are in:

- How much intelligence they place in the agents;
- The nature of the rules (or sensors) employed;
- How rules are added or modified; and
- How they consolidate performance and event data for real-time monitoring and historical reporting.

All these issues can affect the functionality of your monitoring and management solution as well as the system and network bandwidth it requires.

11 Tips for Delivering the Highest Level of Application Performance

Based on Heroix's real-world experience with a wide range of organizations, we have identified 11 proven tips to help you evaluate your current monitoring and management approach. These guidelines will help you decide whether to optimize your existing solution or to create or buy a new one. They will also provide insight into what you should look for in a truly effective commercial monitoring and management solution.



1

Take a unified, single-console approach to managing all your distributed applications.

No matter how large or complex your infrastructure, you need one centralized and authoritative view of all its components. These include:


- Enterprise and e-Business applications
- Database servers
- Messaging applications and servers
- Web servers
- Operating systems
- Infrastructure hardware
- Tools and utilities

The single-console approach provides the all-important view of your complete, end-to-end IT architecture. Combined with intelligent agents at each infrastructure component, it enables you to monitor and manage not just individual components and applications but the *interaction between them*—often a primary source of system problems.

At the same time, single-console monitoring and management makes it easy to give the different support groups within your department access to the information they need. With intelligent filtering at the component level, any system on the network can forward performance and event data to any other location. For example, e-mail issues can be filtered directly to the messaging group, database performance problems to DBAs, e-Business server matters to Web administrators, and so forth.

This means that specialists can spot problems faster, respond to end-user queries with solid information, and get all the diagnostic data they need to take the right course of action. Plus, with information on application status shared at all levels of the IT staff, more resources and intelligence can be brought to bear on resolving system-wide problems.

Centralized View



Compare with the decentralized, multi-console approach, the single-console model provides a true forest-and-trees view of your distributed applications: you can monitor both individual components as well as your entire application infrastructure. It gives you the flexibility to tailor your monitoring approach to your infrastructure and your IT resources. At the same time, the single-console design lets you consolidate data and event reports in a central repository. This ensures that senior IT managers are fully informed at all times on current conditions across the infrastructure and can monitor, in real time, the corrective actions taken.

2

Multiplatform monitoring and management is a “must have.”

The sheer complexity and diversity of today’s information systems demand multiplatform monitoring and management. This doesn’t just mean that functionality can be extended to a range of different applications, operating systems, hardware platforms, and network protocols. It means that monitoring and management must be totally transparent, with no variation in performance and responsiveness for different components, which include:

- Operating systems: Windows NT/2000/.NET, Unix, Linux, OpenVMS, and Novell Netware
- Network protocols: SNMP, TCP/IP, FTP
- DBMS: Oracle, DB2, SQL Server, Sybase

- Messaging servers: Microsoft Exchange, Lotus Notes/Domino
- Web and e-Commerce servers: Apache, WebLogic, Microsoft IIS
- Hardware platforms from legacy mainframes and client/server systems, as well as any new hardware technologies that emerge.

In addition, you should look at how your IT department's support activities and how different support groups relate to each other. Many organizations find themselves challenged with managing support groups defined by computing function (e.g., messaging, Web server, database, networking, etc.) and also by platform. Without sufficient coordination, this can reach an unmanageable level of complexity that leads to less timely problem resolution.

Your monitoring and management approach has to mirror the seamless interoperability of your IT infrastructure. Only a unified, streamlined multiplatform approach ensures the same level of event detection and correction, performance tracking, and fast identification of potential problems and events.



Monitor proactively to spot and resolve potential problems before they become real ones.

Effective application monitoring, like the best medicine, is preventive, not just therapeutic. You need to monitor proactively for changes in traffic levels, processing volume, or any other condition that could lead to interruptions of application availability or to lower processing performance. This includes monitoring disk space, network traffic, processing bottlenecks, application errors, print server issues, and more.

The use of intelligent, distributed agents is the key to robust proactive monitoring. Ideally, these agents should be located at each component point across your infrastructure and as close to the monitored components as possible. With a greater degree of monitoring granularity, you can more quickly pinpoint the causes of potential problems and take corrective action.

Proactive monitoring doesn't just help you avoid service disruptions, it enables you to leverage your staff's limited time and resources. Call it effective time management—by heading off problems before they happen, your staff spends less time putting out fires and can focus their energies on more strategic, future-oriented systems issues.

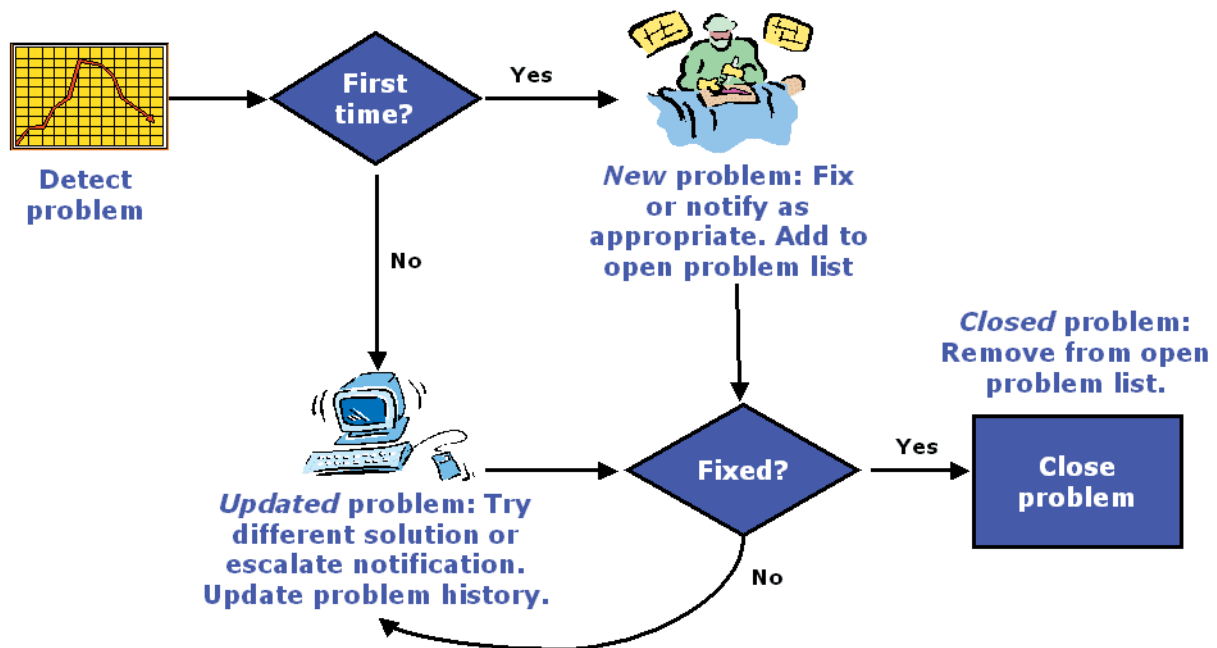
When it comes to corrective action, keep your options open. You may choose to implement monitoring and management that alerts the appropriate person to a potential problem and then requires him or her to take action. Or you can build in rules for escalating problems—issuing wider alerts, triggering phone calls or e-mails to people progressively higher in the IT department hierarchy—until a problem is solved. And last, you can build in automated corrective action that is triggered when pre-defined (or dynamically defined) thresholds are reached or other conditions occur. In all of these cases, you're not limited merely to responding after the fact.

4

Avoid the data deluge—employ closed-loop monitoring to consolidate status reports and simplify the monitoring process.

One of the most common blocks to effective application monitoring—proactive or otherwise—is the sheer volume of performance and event data bombarding support groups. Based on intelligent agent capability, closed-loop monitoring provides a way to simplify and rationalize the management process.

Closed-Loop Monitoring



With this approach, notable performance issues or events can be classified as “new,” “updated,” or “closed.” You can initiate actions (manually or with some level of automation) based on the category. For instance, if an application fails and the status is “new”—and in closed-loop monitoring only the first report of this failure is classified as “new”—you can attempt a restart. If this is an updated report (i.e., a subsequent report of an open issue), you can escalate to the appropriate next step, such as notifying the proper support person.

Closed-loop monitoring has the further advantage of clearing away the visual clutter. Instead of seeing a string of barely distinguishable “events” that merely update an original event, you see one event that is periodically updated until the underlying problem is resolved and the problem report closed.

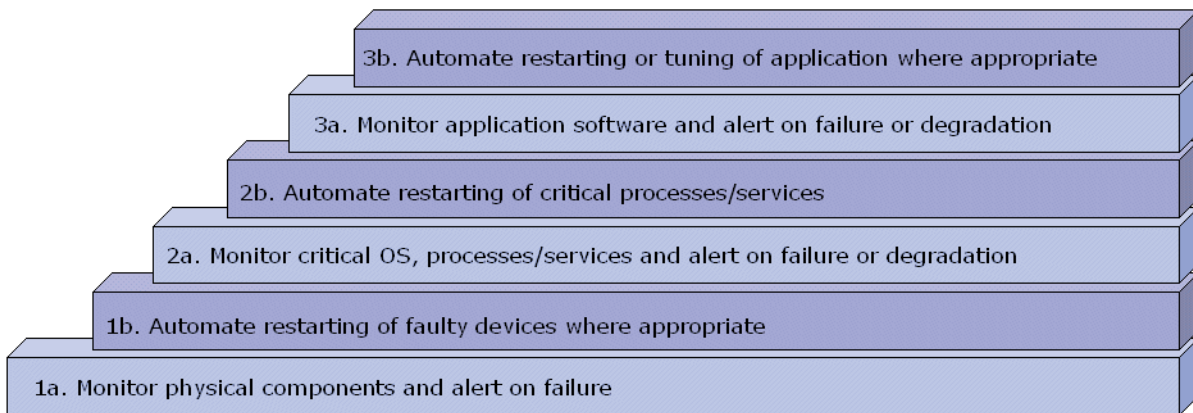
5

Provide for the option of automated corrective action.

To help improve productivity and avoid service interruptions, you should provide the option to use automated problem correction in appropriate situations. Auto-correction can be valuable not only in helping to meet those end-user demands but in maximizing your staff's limited time, technology, and human resources. But many IT managers aren't comfortable with auto-correction across all applications. Some feel it gives them less control and latitude in taking the most effective corrective action.

How extensively should you employ auto-correction? This can be a function, not just of your individual comfort level, but of the nature of your specific IT environment as well as end user or management concerns. If you want to limit auto-correction, you may want to start with "notification only" when an event occurs. Then, as it feels appropriate, you can automate routine maintenance tasks or initiate action when a simple stop and restart may solve the problem. From there, you can extend auto-correction as you see fit.

Levels of Implementation: Corrective Action vs. Notification Only



A staged approach to automated monitoring: First, monitor for problems and send alerts (steps 1a, 2a, 3a). Add automated corrective action based on your comfort level (steps 1b, 2b, 3b).

The more intelligent your solution, the more you may be able to correct automatically and the greater flexibility you will have in customizing your auto-correction responses. For instance, it should let you create and change the monitoring rules (or sensors) to restart a service or application automatically when a failure first occurs, and simply record that action taken. Or you might program it to issue an alert at the same time. That way, you can build in a protocol for escalating problems until they are resolved.



6

Choose a self-configuring monitoring and management solution.

One of the biggest inhibitors to implementing new monitoring and management solutions is the time and cost involved in customizing them to support your particular IT infrastructure—or even just getting them up and running. This can eat up IT staff time and/or involve expensive outside consultants.

To avoid these problems and speed deployment, implement a monitoring and management system that features dynamic self-configuration. With embedded intelligence at the agent level, these solutions automatically perform many basic configuration tasks such as identifying which applications and underlying disks and file systems to monitor and which services to check.

In short, you want a self-configuring monitoring and management solution that you can implement essentially “out of the box,” without additional code writing, scripting, or outside support. Equally important, you should get a system that dynamically re-adapts as conditions change, whether these changes entail variations in usage and traffic levels, the addition (or elimination) of applications or objects, or resizing your entire infrastructure.



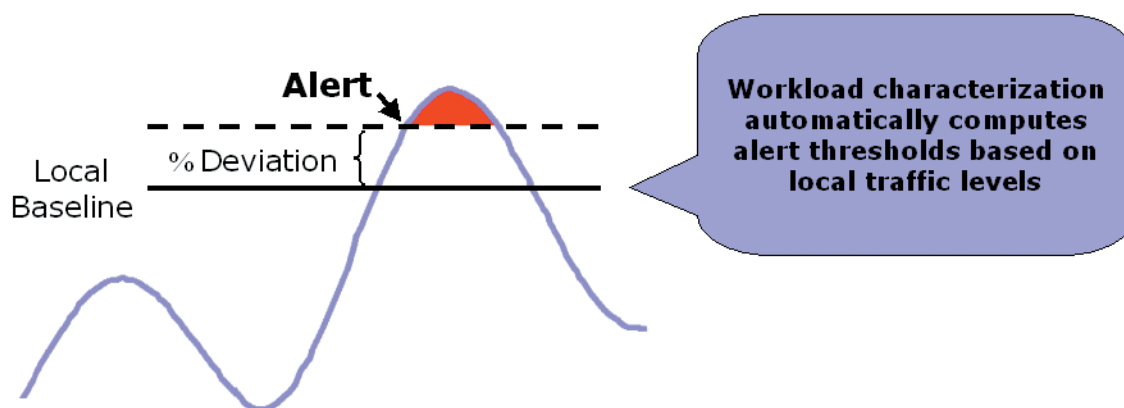
7

Don't do all the work yourself. Let dynamic workload characterization adjust monitoring thresholds automatically.

One of the greatest challenges in application monitoring across the full IT infrastructure is setting and maintaining the right thresholds—for processing loads, disk space, network traffic, output to peripherals, and so forth. One size does not fit all.

What constitutes normal traffic or processing levels for one department or application may be very different for other systems or user groups, and at different times throughout the year. For instance, the e-mail server for a sales and marketing department may carry much more volume than a similar server for an engineering group. As a result, you would need a different, higher threshold on the sales and marketing server before an alert for corrective action is triggered. Conversely, the thresholds on the engineering e-mail server should be lower to avoid letting problems go unnoticed.

Workload Characterization



Trying to maintain the correct threshold levels manually would be, if not impossible, very time-consuming, and because traffic levels change constantly, it would be an on-going process. One of the advantages of a monitoring solution with agent-level intelligence is that it can do the job for you. Through workload characterization, the intelligent agent dynamically discovers new standard traffic and processing levels, and adapts its monitoring thresholds accordingly. This not only saves valuable staff time and energy, it keeps your application monitoring and management at peak performance.


8

Make sure you can easily, seamlessly monitor proprietary as well as packaged applications.

The old ideal of a homogeneous, all-proprietary or all-standards-based application infrastructure is long gone. The typical IT system today is a combination—sometimes a hodge-podge—of proprietary, home-grown applications, and packaged, standards-based software solutions. Many of these have been customized and amended extensively over time.

You need a management solution that adapts to and monitors both kinds of applications with equal ease and effectiveness. Further, it should monitor not just the applications themselves but also the interaction between those applications. For example, if you have a Web-based application that draws upon a back-end database, you should be able to detect problems not just with the front-end system, but with the performance between the two systems.

To do this, choose a management approach that lets you create custom rules and modify existing rules to support proprietary (including legacy) applications and site-specific conditions. With a graphical, task-driven customization interface, you should be able to adapt monitoring parameters—define and combine schedules, thresholds, and corrective actions—without any additional scripting or code writing.



9 Ensure easy, complete scalability and extensibility of your monitoring and management functions as your IT infrastructure changes.

It's a truism that information infrastructures and applications are constantly changing in response to new business conditions: companies grow, contract, merge, and reorganize. Their IT systems have to respond accordingly, enabling the support group to add or drop end users, implement new applications or upgrades, and more.

To keep pace with these changes, you need a monitoring and management solution that is easily scalable and extensible. It should require minimal hands-on effort on your part, in the form of programming and customization, to extend monitoring to new applications, or to enhance your existing monitoring and management capabilities. As discussed above, this begins with automated reconfiguration and application detection to ensure that your monitors can dynamically "recognize" new applications, users, and parameters. Furthermore, you should be able to tailor, scale, and extend your application management in hours rather than days or weeks.


To achieve this speed and agility, you need an interface and underlying architecture that enable you to define changes on an enterprise level and then efficiently communicate those changes to the intelligent agents at the local level. For example, if you want to change the way you respond to a particular problem with a supply chain application, you should only need to make the change once, and then have the change replicated for you to all affected application servers. As a result, instead of having to perform hundreds of operations to update as many servers, you can update all the servers with just one operation.



10 Get the reports and documentation you need to show fulfillment of Service Level Agreements—whether formal or informal.

It isn't enough just to meet your end users' system performance expectations. You need the documentation to prove you've done it. In some cases, documentation and performance reports and graphs are mandated to prove that the company has fulfilled its formal performance commitments to clients. Even when there are no explicit service level agreements, it's wise to set *de facto* performance standards, whether they're formal agreements with line-of-business managers or informal goals only shared within the IT group.

In either case, you need an effective reporting and graphing mechanism to document compliance with these expectations. These reports and graphs should track performance across your entire IT infrastructure and offer clear, easy-to-read data, including graphical display for greater clarity. Further, the reports should be generated and distributed automatically (via e-mail or your intranet), so you don't spend hours extracting and compiling data from diverse sources, then printing and faxing them to recipients.



Beyond showing compliance, detailed performance reports give you a valuable perspective on your own application performance over time. They let you spot ongoing trends, so you can more accurately predict and plan for future IT needs.

11

Ensure 100 percent fault-tolerant monitoring and management through real-time redundancy.

When a component fails and doesn't self-correct, you need to track down and rectify the problem as quickly as possible. You need to identify the applications and locations involved as well as the conditions and thresholds immediately prior to the outage. The only way to get that data is through a fault-tolerant monitoring system, one that records those conditions no matter what.

For that, you need real-time redundancy in the form of primary and backup consoles. If you have intelligent agents across your distributed applications, these agents should continuously feed performance and event data to both the primary and the backup console. So, if the primary console fails, the data continues to flow without interruption to the backup console.

Conclusion

So, how do you deliver the application performance end users want? It's clear that there's no "one size fits all" answer. In today's diverse, rapidly evolving computing infrastructures, IT managers have to be more creative than ever to get the most out of their limited IT resources. It's often a balancing act between commitments to end users and the limitations of IT budgets, computing resources, and staff time. One key variable the IT manager can control is the effectiveness of the monitoring and management solution.

The tips discussed here will give you a clear picture of options available to you—and some of the issues you should consider— in selecting or refining your monitoring and management approach. Further, they will help you gain greater control and flexibility over your monitoring process. Nothing can ensure 100 percent availability and up time, but these tips offer solid guidelines for reaching new levels of performance and for keeping pace with the ever-changing IT infrastructure

About Heroix Corporation

Heroix delivers award-winning software products that help organizations optimize the availability and performance of applications, systems, and IT infrastructure running across multiple platforms. Today, with a presence in over 40 countries, Heroix has delivered performance management solutions to over 1,000 organizations worldwide, including AMD, CSX Corporation, Dollar Rent A Car, GlaxoSmithKline, Hasbro, Honeywell, Kaiser Permanente, Mary Kay Inc., Motorola, Inc., and Sprint.

Clients rely on Heroix's software to detect application and system IT problems, report and graph performance data, track service levels, and centrally monitor and manage wide-ranging technologies including packaged and custom applications, database systems, messaging platforms, Web servers, operating systems, and infrastructure components. To learn more, visit www.heroix.com.

Heroix Corporation believes that the information in this document is accurate as of its publication date; such information is subject to change without notice. Heroix is not responsible for any inadvertent errors.

Heroix, the Heroix logo, and Heroix eQ are trademarks of Heroix Corporation. All other trademarks are property of their respective owners.