

Better Incident Management Through Application Performance Monitoring

Written by
Ken Barrette
Manager, Product Management,
Quest Software, Inc.

© 2010 Quest Software, Inc.
ALL RIGHTS RESERVED.

This document contains proprietary information protected by copyright. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose without the written permission of Quest Software, Inc. (“Quest”).

The information in this document is provided in connection with Quest products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest products. EXCEPT AS SET FORTH IN QUEST’S TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software World Headquarters
LEGAL Dept
5 Polaris Way
Aliso Viejo, CA 92656
www.quest.com
email: **legal@quest.com**

Refer to our Web site for regional and international office information.

Trademarks

Quest, Quest Software, the Quest Software logo, AccessManager, ActiveRoles, Aelita, Akonix, AppAssure, Benchmark Factory, Big Brother, BridgeAccess, BridgeAutoEscalate, BridgeSearch, BridgeTrak, BusinessInsight, ChangeAuditor, ChangeManager, Defender, DeployDirector, Desktop Authority, DirectoryAnalyzer, DirectoryTroubleshooter, DS Analyzer, DS Expert, Foglight, GPOAdmin, Help Desk Authority, Imceda, IntelliProfile, InTrust, Invirtus, iToken, IWatch, JClass, Jint, JProbe, LeccoTech, LiteSpeed, LiveReorg, LogAdmin, MessageStats, Monosphere, MultSess, NBSpool, NetBase, NetControl, Npulse, NetPro, PassGo, PerformaSure, Point,Click,Done!, PowerGUI, Quest Central, Quest vToolkit, Quest vWorkSpace, ReportAdmin, RestoreAdmin, ScriptLogic, Security Lifecycle Map, SelfServiceAdmin, SharePlex, Sitraka, SmartAlarm, Spotlight, SQL Navigator, SQL Watch, SQLab, Stat, StealthCollect, Storage Horizon, Tag and Follow, Toad, T.O.A.D., Toad World, vAutomator, vControl, vConverter, vFoglight, vOptimizer, vRanger, Vintela, Virtual DBA, VizionCore, Vizioncore vAutomation Suite, Vizioncore vBackup, Vizioncore vEssentials, Vizioncore vMigrator, Vizioncore vReplicator, WebDefender, Webthority, Xaffire, and XRT are trademarks and registered trademarks of Quest Software, Inc in the United States of America and other countries. Other trademarks and registered trademarks used in this guide are property of their respective owners.

Updated—October 2009

Contents

- Introduction: The Need for Business and IT Alignment4
- Step 1: Change Your Perspective5
- Step 2: The People7
- Step 3: Better Process8
 - Prioritization.....8
 - Finding the Source9
 - Escalation.....9
- Step 4: The Right Tools Make a Difference11
- Step 5: Don't Forget to Measure12
- Conclusion: Start Somewhere.....13
- About the Author14
- About Quest Software, Inc.15

Introduction: The Need for Business and IT Alignment

With each passing day IT, becomes a more integral element of operational efficiency, product innovation, customer service, and other avenues of competitive differentiation. Naturally it follows that unplanned application outages or even just slow performance have a direct effect on the business in terms of higher costs, lower revenues or lost customers.

With the complexity of applications today, problems like these are bound to happen. Does the IT organization understand business priorities well enough to focus on the most important ones? Does it have the right support processes and effective tools to minimize the business impact of application-related incidents? There are many areas of IT to evaluate for changes to better support the needs of the business, but we need a good place to start.

Improving the mean-time-to-repair (MTTR) of incidents is a tangible way for the IT organization to increase its value to the business. What follows is a recommended set of practices covering people, processes, and tools that will enable an IT organization to prioritize incidents based on business impact and more quickly restore services to operation.

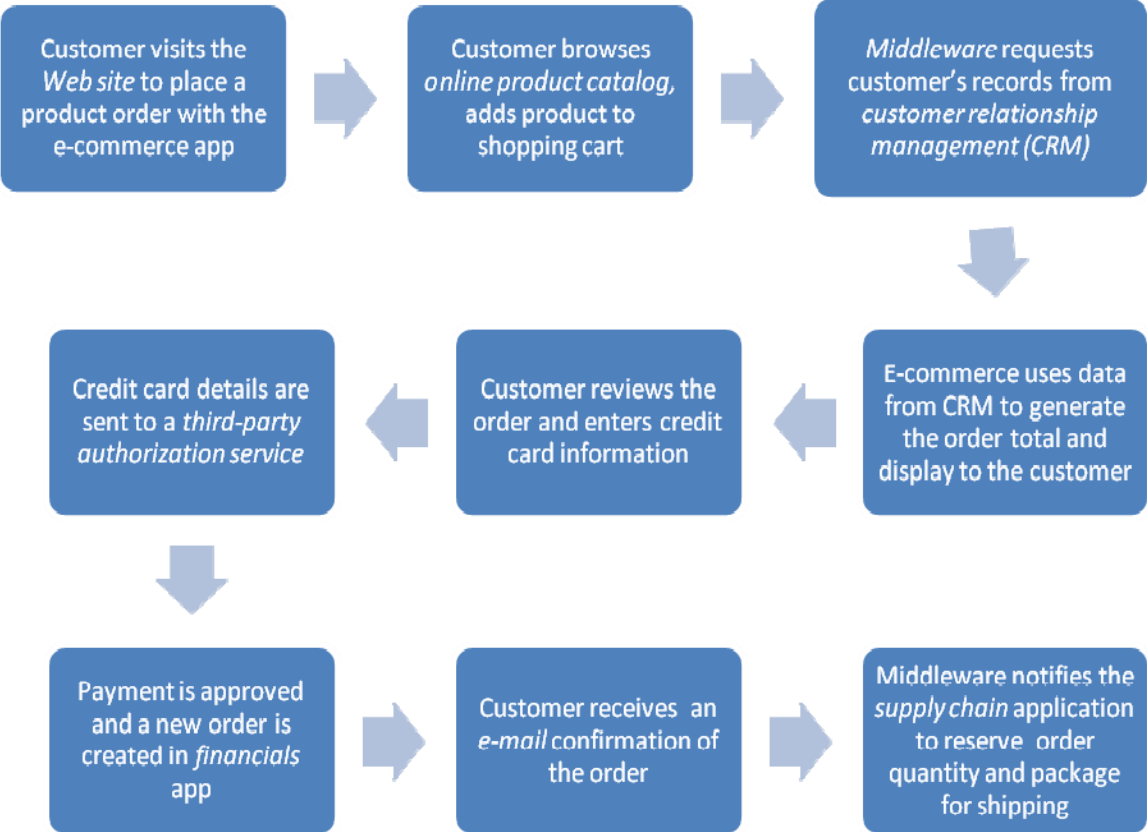
Step 1: Change Your Perspective

Against the backdrop of increased dependency between business and IT, we should reconsider the way we manage IT resources to be more aligned with business services. There is a tendency for IT personnel to manage technology in silos. We manage server and network resources with broad, device-centric tools. We manage application software resources with platform-specific tools. Rarely do we manage hardware and application software as the interdependent collection they really are.

This disconnect becomes more of a problem because neither the hardware tools nor the application platform tools provide a perspective that is naturally aligned to the way the business operates. The user of a business service doesn't care which server runs the software. A transaction may traverse multiple applications, middleware or even third-party services as it executes. The shift towards virtualization of IT infrastructure further compounds the problem since the relationship between infrastructure and the applications they run is more dynamic than ever.

Whether the service in question is e-mail, telephony, accounts payable, or a Web store, the goal is the same—provide visibility into how the service is functioning in its entirety. By continuing to manage silos, IT and business alignment will not be achieved. There will be barriers to collaboration that cause delays in the resolution of any incidents, resulting in excessive downtime.

Consider the following example of an online order processing service:



Even in this relatively simple example, there are potentially nine different applications involved. And that is just the software. The number of servers and network connections involved could easily number in the dozens. In order for IT to support this service, it needs visibility into all of these interdependent components. However, a true service perspective is more than just a cosmetic grouping of the servers, software and devices in a single console. The right perspective provides business context and prioritizes the incidents with the biggest business impact. It shows:

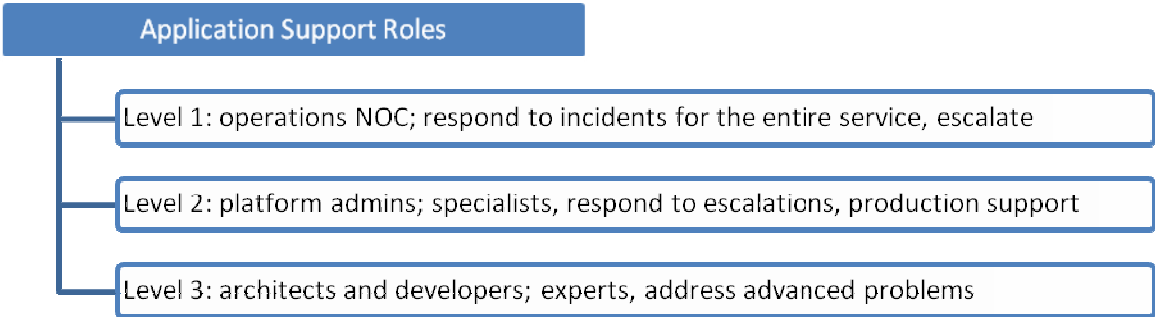
- Dependencies between components and other services
- Performance of individual components
- Success in achieving service level objectives
- Status of important transactions
- Quality of the user experience

This kind of holistic view is the necessary first step toward aligning IT with the business around a common understanding of a service. An incident in any one of the service components can be seen in context of the other components.

Step 2: The People

With the ability to view IT resources as services in the appropriate business context, the next step toward more service-oriented IT management is to establish a support structure to easily identify the business-critical incidents and resolve them as quickly as possible. Understanding which ones are the most important requires IT staffers who not only know the bits and bytes of the technology, but also appreciate the business purpose of the service they support. For example, if the accounts payable service has an outage at the end of the month when the company needs to close the books, that is a bigger problem than if that outage was to occur in the middle of the month.

All these attributes don't need to be present in a single person, and in fact, that may not be reasonable to expect. Even today, IT organizations typically have a multi-level team structure with different and increasingly specialized technical skills at each level. This is the start of a good model to create a more service-aligned support team. The specific names for these roles may vary from one organization to the next, but they typically follow a pattern along these lines:



With their position at the front line, Level 1 support staffers play a critical role. They work directly with the business or customer, gather as much information as possible about an incident, and find a potential resolution or workaround. They certainly don't need an MBA, but they will need some basic understanding of the business functions carried out by the services they support. This contextual knowledge will help them provide better support to their customers compared to just knowing there is an alert on a particular server or application.


While the team structure is set up to handle escalations, too much escalation is a sign of problems with the team. This multi-level support structure should not be a strict hierarchy where activity only flows in one direction. Because Level 2 and Level 3 staffers have a combination of reactive and proactive responsibilities, they cannot attend to the longer range service improvement activities if escalations from Level 1 become too frequent. These more specialized team members can benefit by transferring knowledge to Level 1 staffers as well as creating tools that automate common tasks and resolve more incidents on their own. Enabling Level 1 staffers to resolve more incidents without escalation is an important goal in improving the value of IT to the business. In order to be high performing, the teams need to realize that their relationship is an interdependent, mutually beneficial one.

Step 3: Better Process

A better aligned IT staff with clearly defined support responsibilities and a shared commitment to the customer can certainly improve the quality of the IT service, but there are additional gains to be had in defining an effective process for prioritizing, researching and escalating incidents. While the technologies that support different services will likely vary greatly from one service to the next, that doesn't mean the way that incidents are handled has to be reinvented for each and every service. Consistency of process can lead to better efficiency and enable the staff to scale up to higher levels of service.

In the best-case scenario, an alert is related to an incident that has a known workaround and the problem can be dealt with right away without unnecessary escalation. Realistically, the complex nature of the typical IT infrastructure often results in new incidents that have no known resolution. Incidents of this sort will need to be escalated to Level 2 and in some cases to Level 3.

The key parts of the incident handling process are:



Prioritization

Not all alerts about incidents are to be dealt with equally. How can a Level 1 support staffer determine the right ones to focus on? When a support team is responsible for multiple business services, determining incident priority is not as simple as looking at a red, yellow or green dot on an event monitoring console. There are several factors to consider in determining the priority of an alert:

- Which service is affected and is it important?
- Is there a service level agreement (SLA) violation?
- How many end users are impacted?
- Is the problem intermittent or persistent?

Service importance: an IT support team will likely be responsible for multiple services. These services will have different priority weights depending on the nature of the business, whether they are customer-facing vs. internally used, or even on the time of day/week/month/year.

SLA violation: the SLA defines performance, availability and response time criteria of the service to which the business and IT have agreed. Considering service level impact is the first step in determining relative priorities. An SLA should be tailored to the nature of the service and degree of criticality of its business purpose. With a Web-based catalog service, if a Web page rendering slowdown is not yet affecting the SLA, it would get a lower priority than an alert from the order entry service where the SLA was being violated.

Impact on end users: another key measure of service health is the experience of its end users. In the Web-based catalog example, if slow Web page rendering is affecting three percent of end users, Level 1 support would prioritize that incident differently than if it was affecting 30 percent of end users.

Intermittent or persistent: if an incident occurs once, does it deserve the same priority as an incident that occurs repeatedly over the same period of time? Depending on the previous prioritization factors, even a single incident may be a high priority, but in the abstract an incident that sends the service into a critical state for an extended period of time is likely to be a higher priority than the isolated one.

Prioritization guidelines should be developed with input from all levels of support as well as the business users. Coming to agreement on what constitutes a high-priority incident will minimize the chances for miscommunication and properly set expectations.

Finding the Source

The next part of Level 1 responsibility is to narrow down the source of the alert. In a multi-tier, distributed service this may not be simple. A single alert about degraded performance of user experience may be the result of multiple problems in different parts of the infrastructure. It is also essential to consider what the user saw at the time of the incident.

In an example such as an online ordering service, if a Web page displays 50 percent slower than it did in the previous hour, what is the cause of that slowdown? Is it network latency? Is it due to increased user traffic and load on the Web server? Is it the result of database resource contention? Is the virtualization hypervisor starving the application server of resources? With distributed applications, there could be a variety of causes.

Without narrowing down the source of the slowdown, a fire drill among the administrators for the database, the Web site, the application, systems, and the network could result. The other possibility is that a “pass-the-buck” situation could arise among all the parties where each individual thinks the problem is in someone else’s area of responsibility. For effective incident triage, it’s not enough to monitor the different tiers in isolation. An end-to-end view of the entire service improves the chances of a fast resolution.

Monitoring user activity provides important insight into the source of an incident. It can be monitored passively, actively, or through synthetic transactions. Activity can even be recorded to capture what the users were doing at the time of the error or slowdown and what they saw on the screen. Having this complete range of user experience data is the ideal. Only this comprehensive user experience data provides the full set of clues as to whether the source of the incident is with client/browser, back-end infrastructure or even the network.

If Level 1 staffers can narrow down the source of the alert based on looking at a variety of data points in context with one another, they can take the next step and try to find a resolution for the incident. Sources for this information may be stored in a searchable knowledgebase that the team maintains, or better yet, from intelligent messages that provide suggestions right in context of the alert. By design, the Level 1 staff doesn’t have deep experience with specific platforms, but access to established knowledge and other shared information can expedite restoration of services and reduce the number of escalations.

Escalation

If Level 1 staffers are unable to restore the health of the service, they need to escalate incidents to their counterparts on the Level 2 team. As part of the investigation, Level 1 staffers should annotate or document their findings in some kind of collaborative format such as a service desk ticket, change request, e-mail, or incident report.

For example, take this scenario where an application running on a virtualized server begins to experience a slowdown:

Level 1 staffers do the necessary legwork to gather the details about the specific transaction or request that is executing slowly. They identify:

- The name of the virtual machine it’s running on
- The details of the error message explaining that it’s waiting for CPU cycles

Armed with this information, the Level 2 virtualization administrator has a head start on resolving the issue. Options include:

- Starting up a new virtual machine to add more capacity immediately
- Adjusting the configuration of the virtual machine settings so the issue is less likely to recur

In this scenario, the virtualization administrator may be already aware of the original incident via alerts he received directly. However, the virtualization administrator is unlikely to take any action until he is made aware of this as a high priority. The Level 1 support team provides the critical extra context to let team members in other support levels know how and why this issue is important to the business. Additionally, by narrowing down of the source of the incident to the virtualization layer, time is saved by avoiding repetitive research.

Having this multi-level team and an established process for handling incidents makes the coordination between the different teams as smooth and repeatable as possible.

Step 4: The Right Tools Make a Difference

While the right process can—and arguably should—exist independently of any specific tools, the right tools can make the right process easier and faster to perform with greater reliability and effectiveness. Getting people to change behavior and follow a process is difficult enough. If the tools they use to do their job support the process and make it more natural behavior, chances of success will improve. Even if you have a good team structure and an established process, evaluating the tools used by those teams may yield some additional opportunity for improvement.

In the case of incident handling, even though the staffers in the various levels of support have different skills and responsibilities, this does not mean they should use disconnected tools. The right solution will collect all the necessary service health data, and present it in a manner that makes it as easy for Level 1 staffers to prioritize the incident as it is for Level 2 or Level 3 staffers to diagnose the cause. Without this shared but tailored perspective, the best process in the world will be subject to human error and create redundant sources of information that complicate the incident resolution efforts.

The right view for Level 1 staffers has:

- A list of all business services for which they are responsible
- A summary of health related to the SLA and different parts of the technology stack
- Easy access to the see status of end-user quality of experience
- Appropriate details of alerts to provide context for troubleshooting
- Functionality for creating service desk tickets and annotating alerts

The right view for Level 2 and Level 3 staffers has:

- Highlighted alerts to show what is already being worked on
- Platform-specific views of configuration and performance details
- Diagnostic workflows
- Component-level details (e.g., methods, scripts, line-of-code)

The less time IT personnel spend repeating the efforts of their co-workers, the faster service can be restored and business can operate as usual. Faster incident resolution is not only good for the business, but it also enables Level 2 and Level 3 support staff to spend less time doing the more routine data gathering activities and spend more time being proactive—performing preventive maintenance and researching root causes of problems.

Step 5: Don't Forget to Measure

There is a well-known adage that says: "You can't manage what you don't measure."

For incident management, there are several relevant key performance indicators (KPIs). Just like the IT support staff needs a multi-dimensional view of your business services to determine the source of an incident, IT management needs to measure the process results from multiple dimensions to determine if the incident management process is effective or if change is necessary.

Here are a few KPI examples to consider:

- Time required to restore service
- Number of escalated incidents
- Percentage of SLA attainment
- Number of incidents processed
- Ratio of administrators to servers

Whether it's with these KPIs or others, using multiple metrics on process performance enables informed analysis on areas for improvement. The metrics may suggest a need for more training of Level 1 support staff, changes to application code, faster or better hardware, or more network bandwidth.

These KPIs are also invaluable for communicating the value of IT to the business. If the business requires a higher level of service from IT, the data provides a factual basis upon which the potential tradeoffs of that higher cost associated with that higher level of service can be evaluated. As IT moves more and more toward a service model of its own, communicating the value of that service will be essential.

Conclusion: Start Somewhere

Better alignment between business and IT is a worthy goal with which no one would likely disagree. But how do you actually do it? Where do you start?

Changing established behavior is a challenge. Inherently there is inertia with the way things are currently done. People are comfortable with it and resist change. However, the pressure on IT to support the increasing demands of the business without a corresponding increase in costs will only continue to rise. There are ways to change the existing process that benefit both IT and the business, but it requires a change to how we manage IT resources.

As long as you have the broader goals in mind to improve quality of business service and establish a real business and IT partnership, you don't need to have the perfect plan. Starting with something and iterating to improve is better than suffering from paralysis of analysis while the day-to-day IT challenges continue to mount. Having the long view is not incompatible with making short-term progress. The key is to focus on areas where progress can be made (or at least determined) sooner rather than later.

Achieving higher availability and performance of critical services may be the most fundamental goal shared between business and IT. Faster resolution of incidents is essential to improving availability and performance. Consequently, making IT incident management more effective is a great way to deliver tangible value to the business.

Incident management has the additional benefit of being daily activity, whether there is a designated process for it or not. By establishing some structure around this activity, measuring it becomes possible and you get relatively quick feedback on what is working and what is not.

You don't have to "boil the ocean" and tackle all critical services at the same time. Start with one or two business services. Determine an initial SLA based on what the business needs, not what IT finds easy to measure. Identify the high-priority user interaction points. Define the support roles, responsibilities, and escalation rules. As you learn what works and what needs to be improved, new services can be brought online.

To paraphrase a famous quote, "perfection is the enemy of progress." Waiting to reorganize how IT resources are managed is not an option for operating a successful business in the 21st century; it's a requirement.

About the Author

Ken Barrette is a 17-year veteran of the software industry. He is currently a manager of product management for Quest Software and has been with the company since 2002. His current responsibilities include determining the strategy for the service management capabilities of Quest's Foglight product line. Before joining Quest, Ken worked in product management and product marketing for several software companies including Serena, Dazel (now part of HP) and MessageOne (now part of Dell). He holds a bachelor's degree in business administration from the University of Connecticut.

About Quest Software, Inc.

Now more than ever, organizations need to work smart and improve efficiency. Quest Software creates and supports smart systems management products—helping our customers solve everyday IT challenges faster and easier. Visit www.quest.com for more information.

Contacting Quest Software

PHONE 800.306.9329 (United States and Canada)

If you are located outside North America, you can find your local office information on our Web site.

E-MAIL sales@quest.com

MAIL Quest Software, Inc.
World Headquarters
5 Polaris Way
Aliso Viejo, CA 92656
USA

WEB SITE www.quest.com

Contacting Quest Support

Quest Support is available to customers who have a trial version of a Quest product or who have purchased a commercial version and have a valid maintenance contract.

Quest Support provides around-the-clock coverage with SupportLink, our Web self-service. Visit SupportLink at <https://support.quest.com>.

SupportLink gives users of Quest Software products the ability to:

- Search Quest's online Knowledgebase
- Download the latest releases, documentation, and patches for Quest products
- Log support cases
- Manage existing support cases

View the Global Support Guide for a detailed explanation of support programs, online services, contact information, and policies and procedures.



5 Polaris Way, Aliso Viejo, CA 92656 | PHONE 800.306.9329 | WEB www.quest.com | E-MAIL sales@quest.com
If you are located outside North America, you can find your local office information on our Web site.

© 2010 Quest Software, Inc.
ALL RIGHTS RESERVED

Quest Software is a registered trademark of Quest Software, Inc. in the U.S.A. and/or other countries. All other trademarks and registered trademarks are property of their respective owners.
WP-BetterIncidentManagThroughApplPerfMonitor-US-AG-20100616