



Best practices for problem determination mean faster problem resolution

Contents

This white paper provides a look at best practices as related to the establishment of a basic, effective problem determination approach for your company.

What if your Web site were suddenly unable to process credit card transactions? And if your customers can't make purchases, your company can't make money. So is it the database? Is it the application server? Maybe it's the operating system? You need answers but don't have a lot of time to find them.

Pinpointing these answers costs not only the time and money associated with downtime, but the time your IT staff spends identifying the cause of the problem as well. And as the problem is often buried deep inside a complex IT environment, it's potentially a significant piece of time your staff could be spending more productively.

This process of determining the root cause of a problem, or isolating why a failure occurred, is known as problem determination. To consistently reduce the time to resolution of system issues, it can be helpful to adopt best practices in this area, defining and adhering to a logical, organized, and well-tested approach to problem determination.

Best practice one: common event logging

Without using best practices for problem determination, it might take hours to simply identify a problem. You have to examine the error log for each component individually. And most likely, each of the resources in your environment has different logging formats.

Take the example of one of Austria's biggest data processing centers. With more than 600 employees serving the country's largest banking institutions, it has an infrastructure that includes 5,000 network devices, 80,000 network interfaces, and 20,000 IP networks—generating one million network events per day.

Now imagine trying to correlate and analyze the millions of events these devices generate—such as SNMP-traps, system logs, and network events—with no solution for filtering out unwanted events, and no means to compare events from one log against events in different formats from other resource logs!

So how did this IBM customer address this first step in the problem determination process?

They began by incorporating log analysis capabilities, provided by IBM, into an existing centralized network logging database solution. These capabilities helped convert firewall and system log data into an industry standard format, the Common Base Event format, for accurate comparisons of data. Using this log analysis capability gave the company visual comparisons, time-based correlation, and extended correlation capabilities across multiple logs and multiple formats.

It's also helpful to be able to sort and sequence events based upon common event elements such as timestamp, severity, process ID, and location. Otherwise, how do you pinpoint the problem? When you can correlate events in real time—in a standard format—helping locate the source of the problem quickly, how much time and money can that save?

Practice benefits:

- *Adoption of an industry standard log format, in this case Common Base Event, as a corporate standard enables consistent, good quality logging across all applications*
- *Use of event correlation capabilities enhances existing tools and processes*

On demand business benefits:

- *Consistency, standardization and efficiency improvements for the problem determination process*
- *Faster IT problem determination for better system availability and happier customers*

Best practice two: effective filtering strategies

Importing multiple log files from distributed sources into a single user interface creates a standardized view, in Common Based Event format, of log file data that can be correlated to speed problem determination. But this ability by itself is just the start. At this stage, you still have the equivalent of a very large fire hose pointed at a very small bucket. Although you standardize problem data into a common format across all your resources or devices, the sheer volume of the data makes correlation a real challenge. To speed problem determination, a filtering strategy is an absolute necessity.

When a problem first occurs, the large number of resources, applications, and log file events related to the problem may be overwhelming. So, your goal is to zero in on the key events that are most likely tied to the problem you are investigating. In a typical scenario, the first step might be to examine only events with a “warning” or “error” severity level in logs of resources suspected to be related to the problem at hand. Once an error is identified, the next step might be to use the creation time of the event (or events) to import events from the various logs between a specific timeframe. From here, you could further reduce your result set by filtering against other event source characteristics like location or process and thread ID.

Imagine how much time – and how many resources – you might need to implement this type of filtering strategy manually. Implementing tools to automate your filtering strategy is key to reducing the time needed to isolate causes of a failure. For example, one IBM customer, an energy company, was able to reduce 500MB of log files down to 2MB in 15 minutes by implementing a filtering strategy, using IBM log analyzer technology to automate the process. This might have taken days, even weeks in some circumstances, if approached manually. The customer implemented a strategy of filtering based on transaction size, where filters were used to identify Web transactions containing over 100,000 bytes of data. This result was further reduced by excluding all transactions that contained static data such as bitmaps and executable files. Based on history, this company’s problems usually occurred when individuals were running certain large transactions over and over again. By automating this filtering strategy, they were able to free up valuable IT resources usually tasked with sorting through log files manually, pulling out events relevant to specific transactions.

Practice benefits:

- *Identify a filtering strategy that helps focus on only the information relevant to your problem*
- *Implementing automation technology that filters all the irrelevant error log data, allowing staff to spend less time finding information and more time working on problems*

Best practice three: capture problem history and make it accessible

The lack of a central repository housing past problems means that when an issue does recur, your best hope to fix the problem may be the person who solved it the last time – if they are still around.

Since you cannot guarantee that the same problem will never occur again, problem determination best practices involve learning from past experiences by insuring that information on past problems is captured and shared with the team. Because people change jobs or leave the company, you can't assume the one expert able to isolate the problem last time will still be around. In addition, you can't always assume your staff will remember what happened. You need to capture all information about a problem, including the indicators leading up to awareness of the problem.

For example, when a problem with a Web application first occurs, a user calls the help desk with a specific HTTP error message on their screen. The help desk may pass this along to be investigated as an application problem when in reality, the problem was caused by another resource in the network – perhaps a database server going down, or a port not available due to a faulty router. Today, some companies collect information on past problems, but they may not keep information relevant for use by the rest of the company. Or the information may be saved in documents that aren't searchable, aren't updated frequently, or aren't distributed automatically.

A searchable symptoms database or repository of known problems, symptoms, and resolutions is needed for effective problem determination. IBM log analyzer tools integrate with a symptom database for fast identification of known problems and related resolutions.

Practice benefits:

- *Quicker response to problems*
- *A single repository for viewing past problem history*
- *Knowledge sharing of problems/resolutions among team members*

Best practice four: deploy standardized data monitoring and gathering

Like most enterprises today, you're faced with supporting a rapidly increasing number of applications running on a distributed architecture. You also face the challenge of integrating a wide range of components and applications from many different vendors. Given the distributed dimension, a simple end-user transaction can involve many different pieces – from the Web browser, to the router, to the mainframe IBM DB2® database, and through the application servers and transaction middleware.

Problem determination goes beyond filtering and provides root cause analysis and resolution. Monitoring and Enterprise Console solutions provide sophisticated, automated problem diagnosis and resolution to improve system performance and reduce support costs. In addition, auto-discovery features allow you to understand the environment and process events appropriately. Standardized monitoring of resources can help you optimize the performance and availability of your entire IT infrastructure. With a single customizable workspace portal, you can proactively manage the end-to-end health and availability of your IT infrastructure – including operating systems, databases and servers – across distributed and host environments. IT staffs can also detect bottlenecks and potential problems in essential system resources, and automatically recover from critical situations to ensure business-critical applications are up and running. By standardizing on a common monitoring tool, or by centralizing control, you can take advantage of historical and real-time reporting, allowing your staff to quickly access the information they need to rapidly identify, diagnose and resolve situations.

Best practices and the IBM Build to Manage Toolkit for Problem Determination

IBM recently implemented a new tool to aid in the establishment of best practices for problem determination. The IBM Build to Manage™ Toolkit for Problem Determination helps catalog problems as they occur within your IT systems – then works as a sort of cheat sheet allowing your operations staff to compare symptoms against past problems and remediations. This saves both time and money, reducing the time to resolution for IT issues and therefore, system downtime. Incorporating such a continually evolving symptoms catalog enhances self-diagnosis capabilities, and creates a growing ability to self-manage over time. This reduction of system downtime and root cause analysis can be a fundamental tool in a successful implementation of best practices for problem determination.

The IBM Build to Manage Toolkit for Problem Determination provides symptom editing tools, tutorial videos and support, and is currently used by such tools as the IBM Log and Trace Analyzer. Symptom catalogs produced by the toolkit are also being incorporated into offerings within the Tivoli®, WebSphere®, and Rational® software portfolios from IBM.

The toolkit software is based on a standard submitted by IBM to the OASIS (Organization for the Advancement of Structured Information Standards) board known as the Web Services Distributed Management (WSDM) Event Format (WEF), and IBM's initial implementation of that standard referred to as the Common Base Event.

The bottom line

When a problem occurs, the first priority is to get the application or resources back up and running as quickly as possible. But then you're faced with identifying why it happened in order to prevent another occurrence. In many large corporations, the IT infrastructure contains so many disparate components that this process requires a manual correlation effort, often of Herculean proportions. IT must collect data from different environments, systems and networks and sort through it all for answers. Why not automate the process as much as possible?

Using these problem determination best practices, you gain the ability to screen and pinpoint the key log events in each component of the transaction. Log and trace files from all types of IT components, typically available to only experts on each component, can be easily viewed side by side. IT staffs spend less time decrypting error logs and more time finding the root cause. Furthermore, all support staff can access log file content, giving the log files a realistic, usable role in diagnosing application problems and transaction failures.

IBM problem determination technology, embedded in a number of IBM products, along with an industry standard for log formats can serve a crucial role in your company's IT roadmap. Beyond merely delivering these technologies, they provide the capabilities needed to evaluate system data for events, alerts and errors, ultimately saving both time and money.

As a result, your staff can focus on higher value tasks that better align IT and business goals – helping grow revenue and contain costs, build responsiveness and agility into the organization for increased competitiveness, or enable every employee to be more effective.

For more information, contact your IBM sales representative or visit:
ibm.com/autonomic/pd



© Copyright IBM Corporation 2006

IBM Corporation
Software Group
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
11-06
All Rights Reserved

DB2, IBM, ibm.com, the IBM logo, Rational, Tivoli and WebSphere are trademarks of International Business Machines Corporation in the United States, other countries or both.

Other company, product or service names may be trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. Offerings are subject to change, extension or withdrawal without notice.

All statements regarding IBM future direction or intent are subject to change or withdrawal without notice and represent goals and objectives only.