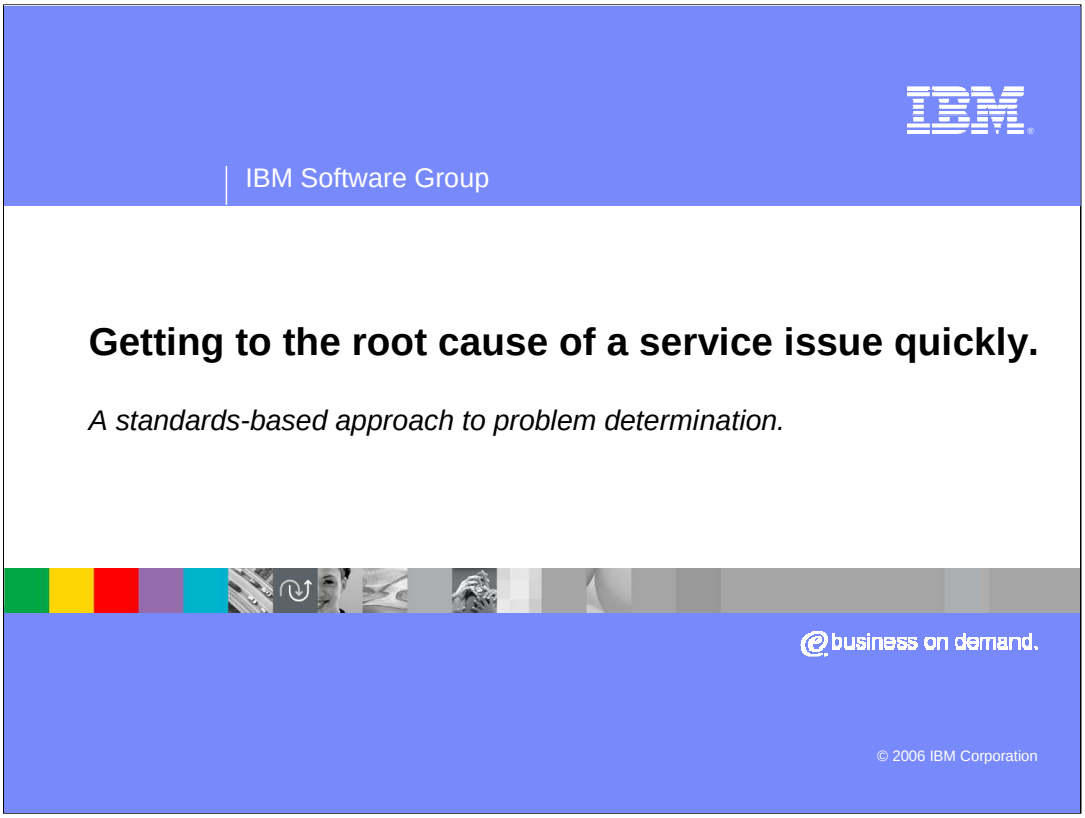

This presentation focuses on a specific challenge many IT organizations face when a system outage or issue occurs. That challenge is problem determination, which entails isolating the root cause of a glitch that negatively impacts availability.

You will be provided with an overview of what IBM is doing to help companies find the root cause of problems quickly, even across vast, complex IT environments.

When a problem occurs, what do you do?

- Priority #1: Recover
  - Get systems back up and running
- Priority #2: Figure out what happened
  - Locate the root cause

This presentation focuses solely on basic problem determination, locating the root cause of problems so that they do not recur.

Clearly, when loss of service occurs your initial focus is on recovering as quickly as possible. And as you are probably aware, IBM offers a comprehensive portfolio of products and services to help you meet required service levels.

But once system recovery has taken place, you must shift your focus to root cause analysis so that the same problems do not recur in the future.

The reason for concentrating on root cause analysis is because problem determination processes commonly take considerable time and resources, which could be better spent on bigger picture priorities, such as implementing new business processes that provide value-add to your customers and grow your business.

So why is problem determination itself such a problem?

The challenge

- IT skills are device or application specific
- Administrators search for answers manually
- Organization silos can create inconsistency

You will probably find the following situation quite familiar. A call comes in from the help desk that customers are complaining about a lag in response time or a service outage. The website is down. Customers cannot access their account information or purchase products, or they cannot make account changes. Whatever the application, it is not working properly. And what happens within your IT department?

Dozens of skilled IT administrators, including your Windows® experts, your database experts, your integration experts, your application experts, and your storage experts start printing out and scanning logs. They begin reviewing infrastructure monitors and other tools, many of which are focused on individual products only. And while some of these tools might identify the error, they still might not be able to isolate the root cause without the assistance of skilled administrators.

Next, your experts gather in a conference room for meeting after meeting, to correlate the data and figure out exactly what happened. Sometimes they only need the logs from that day. Other times, they need the log files for several weeks or months. Often, some finger pointing occurs, a process commonly referred to as “blame storming.” Isolating the root cause, the single issue that caused the cascade of events to finally put service delivery in jeopardy, can take many hours, sometimes days, weeks or even months.

The reason is that problem determination typically is a highly manual process based on a siloed approach. Almost all applications and hardware devices generate log files that capture information about what is going on within that component. And each vendor typically uses different log file formats to report on the health of their device or application. Even products from the same vendor sometimes use different log file formats due to vendor acquisitions and mergers.

Now this is beginning to change, and we will discuss more on that in a minute. But for most companies today, this inconsistent method for event recording, filtering and reporting requires highly skilled IT administrators who are trained in each specific technology component to make sense of it all. And this can be challenging because knowledge and skills are often specific to each administrator. For example, while all your Unix™ administrators might have the same training, each one will likely have his or her own expertise, based on historical experience, for how they solved a particular problem in the past. If the Unix expert who resolved an outage several months ago is now on vacation or has left the company, and that problem recurs, another Unix administrator might have to go through the same discovery and resolution process as his colleague did the first time around.

Complexity is accelerating

- 80% of time is spent on isolating and diagnosing problems[1]
- 60% of availability and performance issues are caused by IT changes[1]
- 80% of development funds are spent identifying and fixing defects[2]

[1]Based on IBM customer engagements
[2] Based on National Institute of Standards and Technology study

When a problem does occur, many of our customers report that 80 percent of IT operations time is spent on isolating and diagnosing that problem. Additionally, according to the National Institute of Standards and Technology, 80 precent of development funds are spent identifying and fixing defects. Given the huge drain on resources, your staff is limited in its ability to deploy new applications and services that the business needs.

And because of the intricate dependencies across the environment, often changes in one area, such as a seemingly simple modification to a database table, can unintentionally cause problems in another area. For many of our customers, more than 60 percent of availability and performance issues are caused by some IT change.

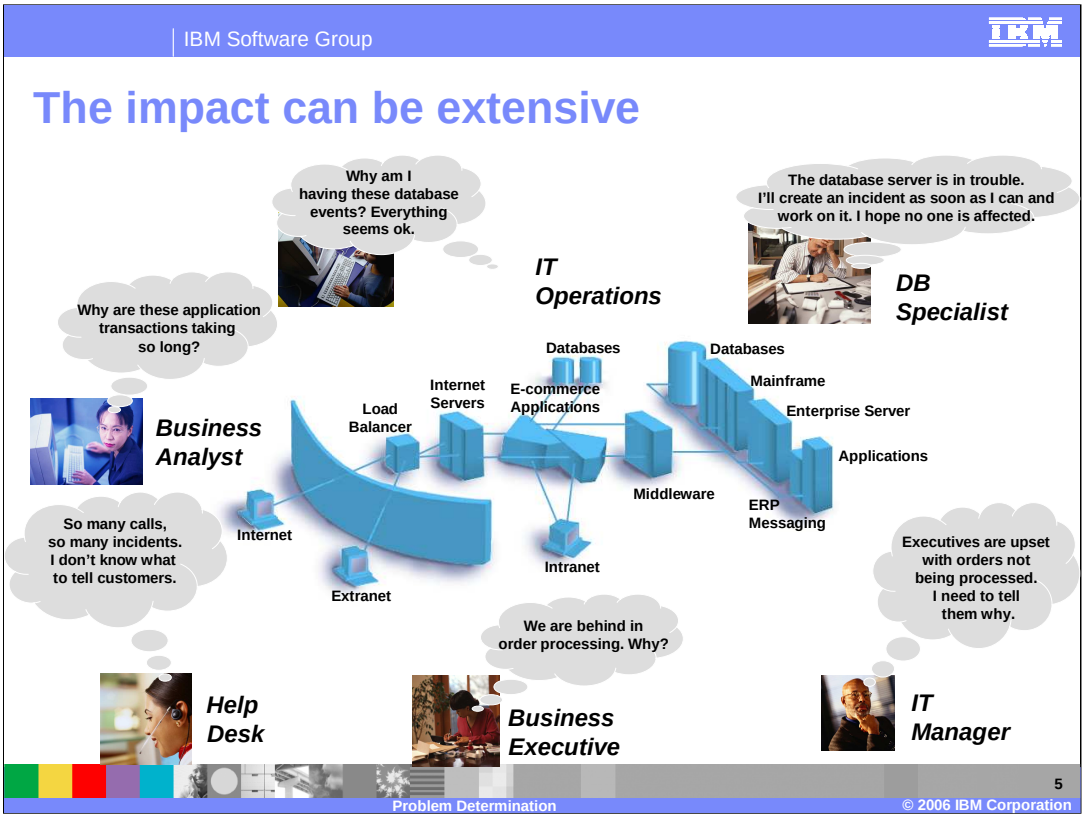

What happens if your staff cannot identify the root cause of a system outage quickly and service delivery problems keep popping up? You will hear about it.

You are likely to receive calls from business executives across the company asking why services are slow or unavailable, and demanding to know when the problem will be resolved.

You are likely to receive calls from your help desk team asking what they should tell customers. And all the while your IT staff is working hard to find answers.

Yet, as your company leverages service-oriented architecture to deploy more complex composite applications, finding the root cause of a problem will only become more difficult.

Stop and think about it. Fifteen years ago a banking service might run on a single mainframe, be accessed by customers using one type of client, such as an ATM, and share information over a private network using a single protocol. Today, that same service is supported by dozens of servers, network protocols, firewalls, security applications, banking applications, and databases.

This creates a huge drain on resources and can result in a greater number of IT administrators on hand with fewer hours to focus on strategic business opportunities, such as building greater responsiveness into your organization, helping improve collaboration across your company, or enabling employees to gain greater insight from existing information.

Whatever your business priorities, think about what your IT administrators could be doing if they were not sitting in a conference room trying to find that needle in a haystack.

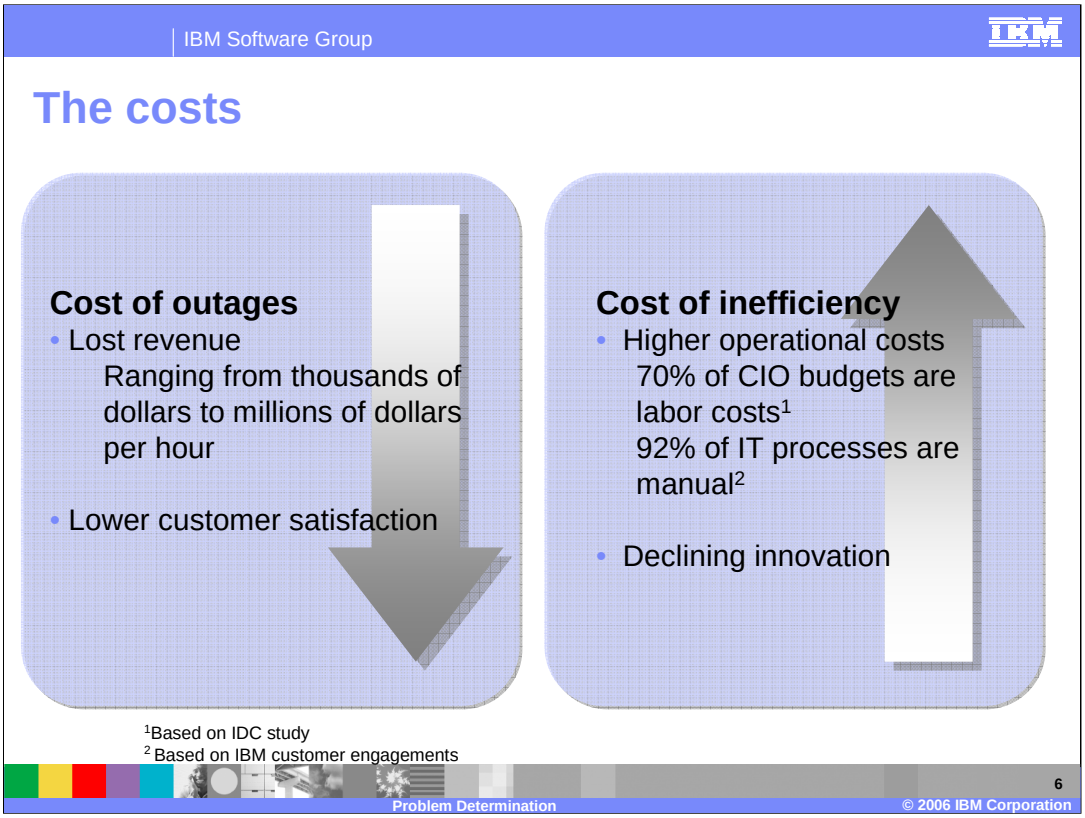

Most people focus first on the cost of outages...and it is a significant concern. The longer it takes to resolve an outage, the higher the cost and the lower customer satisfaction. Numerous studies have shown that the cost of an outage can range from thousands of dollars to millions of dollars per hour, depending on a variety of factors, such as your industry and even the time of day of the outage. In fact, industry analyst EMA found in its research that the cost to one particular company if it lost all of its online systems could reach $2.3 billion per minute.

However, while outages are expensive and highly visible, they might not represent the most significant problem for your organization.

Companies today use a variety of architectural approaches and IT management tools to minimize the impact of IT problems on end-users. But what these processes do not address is the time drain on your IT staff, because the problem still must be isolated, analyzed, resolved and then evaluated to prevent it from occurring again.

And the cost of this is significant.

We commissioned an IDC study that found that 70 percent of CIO budgets are labor costs, and it is no wonder with 92 percent of IT processes today performed manually. This means a greater number of IT administrators on hand and fewer hours to focus on strategic business opportunities.

A better way

- Delivery of industry standard event reporting format

The good news is, IBM is helping companies address this challenge.

Let's focus on *one* aspect of our strategy, namely helping companies bring together the disparate pieces and parts by embedding three innovative problem determination technologies within IBM software products.

This will help you reduce the time to identify the cause of unplanned outages, enabling a more effective and efficient use of resources and strengthening your ability to meet service-level agreements.

The first problem determination technology discussed is the use of the industry standard event reporting format.

Log files from different products often all have their own reporting formats. As co-chair of the Technical Committee formed in the Organization for the Advancement of Structured Information Standards (OASIS) IBM worked with industry participants to create a common language and format for describing problems, allowing faster integration and analysis of events from multiple IT resources comprising complex systems.

OASIS is an international consortium founded in 1993 and drives many of today's e-business standards.

The result of this effort has been the OASIS Web Services Distributed Management Event Format, commonly called the WEF standard. This standard is being adopted by vendors across the industry.

IBM's initial implementation of WEF is called Common Base Events. IBM has taken the base WEF standard and enhanced it to provide administrators with additional information about the resource. Currently, IBM has many hardware and software products that support its initial implementation of WEF, and that number is growing every day. And as OASIS continues to update and enhance the WEF standard, we will provide a migration path.

A better way

- Delivery of industry standard event reporting format
- Ability to quickly convert logs to standard format

It could take a few years before all products in your environment, including IBM and third party, natively support the WEF standard.

Because of this, adapters that can transform proprietary logs into this standard format are currently bundled within select products. These adapters cover many popular IBM and third party applications, including Oracle e-business suite, BEA Weblogic Server, and many more. Our adapters also cover many platforms, including Windows, Linux, z/OS, i5/OS, AIX®, Oracle and Solaris.

Tools have also been embedded within IBM products that include configuration editors and rules to assist you in converting log formats into Common Base Event format. As a result, if you use a product for which we have not created an out-of-the-box adapter, your staff can still quickly convert the log files into the standard format.

The ability to standardize event reporting is an important capability in streamlining root cause analysis, because it allows your staff to more quickly analyze events from heterogeneous IT resources.

A better way

- Delivery of industry standard event report format
- Ability to quickly convert logs to standard format
- Viewing, analysis and correlation of log files

Additionally, innovative technology is incorporated in our products that enables end-to-end viewing, analysis, and correlation of log files across the heterogeneous environment. For example, many IBM products include a consolidated environment that can analyze the logs and traces produced by various components within the application ecosystem.

This technology provides easy “drill-down” capabilities so your staff can see exactly what is causing each problem. With this capability, your administrators can identify the root cause of the most complex problems in heterogeneous, distributed infrastructures, and do so in a fraction of the time.

This self-managing autonomic technology is called IBM Autonomic Log and Trace Analyzer. In addition to embedding this innovative technology into IBM products such as WebSphere Application Server, IBM has also submitted it to Eclipse as open source to enable everyone to leverage this capability. This has been made available in the Eclipse Test and Performance Tools Platform which was recently released as part of the Callisto release.

For customers using IBM WebSphere software, this environment links to a WebSphere Symptom Database that can match patterns, associated solutions and directives. As a result, your administrators can rapidly identify repetitive problems and see how they were resolved previously.

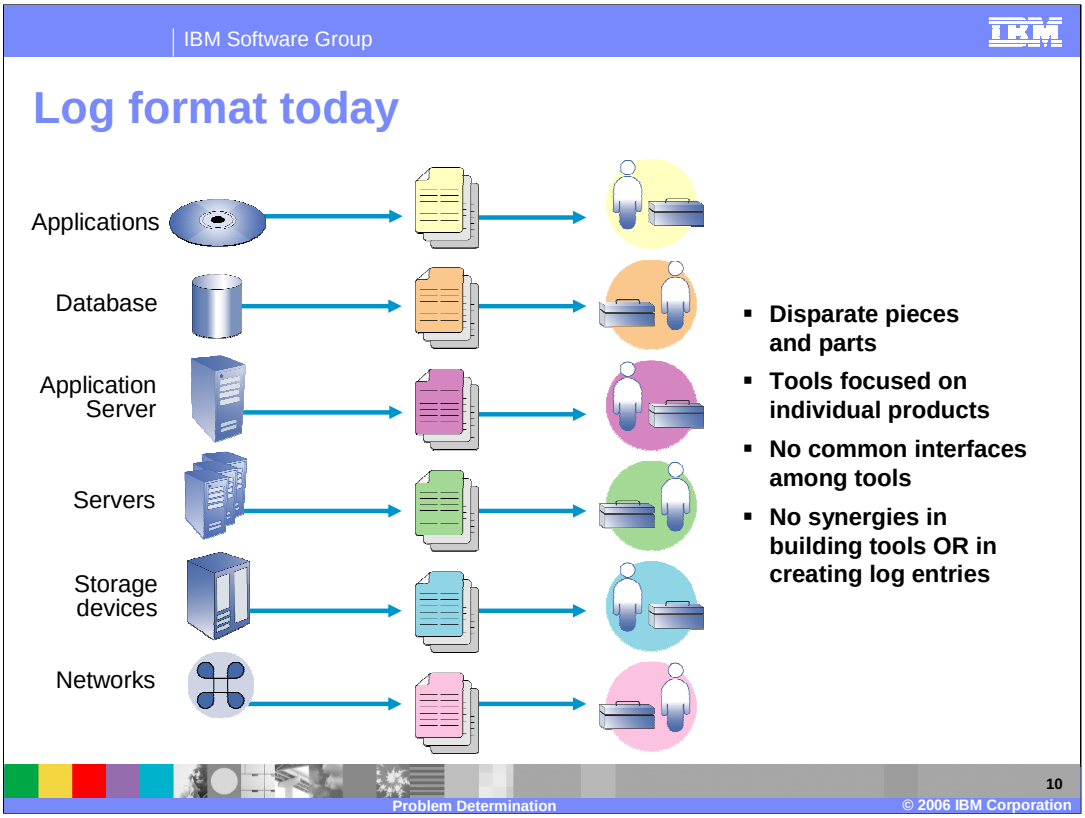

To reiterate, this is the situation today, in which each system produces its own log files which must be manually correlated.

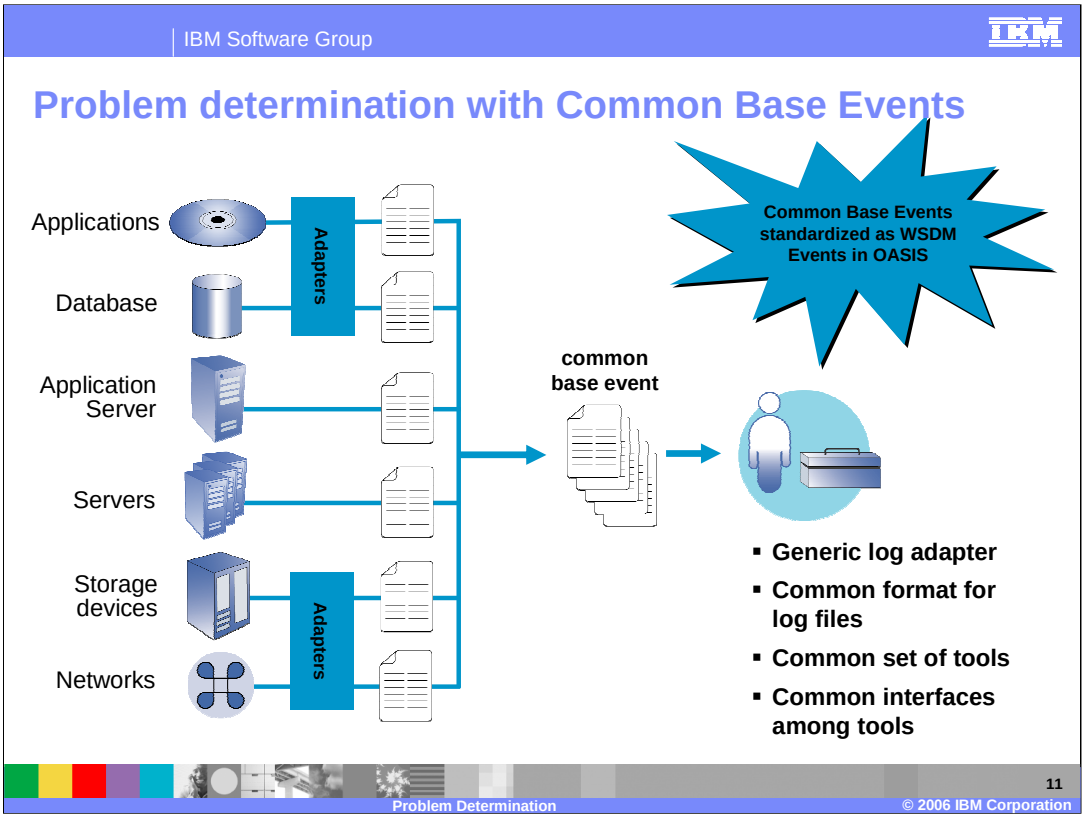

This is the solution made possible by Common Base Events, and made part of the Application Server Toolkit inside WebSphere Application Server.

Now let's take a look at what can be done with this new set of capabilities.

Log and Trace Analyzer

Enables end-to-end viewing, analysis, and correlation of log files across the heterogeneous environment

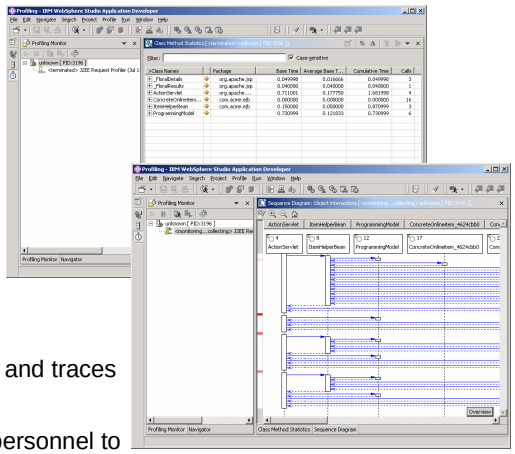

Value:

- Viewing, analysis, and correlation of log files
- Consolidated environment that deals with logs and traces produced by various components
- Easier and faster for developers and support personnel to debug and resolve problems
- Link to WebSphere symptom database available today

The Log and Trace Analyzer is available inside of the Application Server Toolkit component of WebSphere Application Server. It provides the means by which problems can be solved involving not just WebSphere Application Server, but many of the other systems integrated with it, such as operating systems, Web servers, storage, integration middleware, and applications. For a demonstration of the capabilities of this tool, please click on the self-running demonstration in this section.

Symptom Database

- **Used in the analysis of events and error messages that may occur in a log.**
- **XML file of symptoms, string match patterns, associated solutions, and directives.**

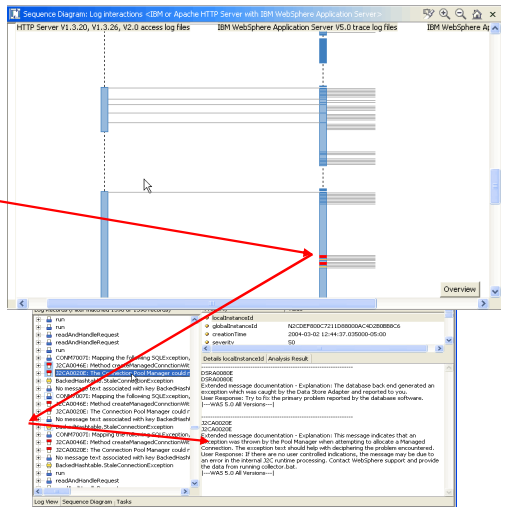

The symptom database is a feature within the Log and Trace Analyzer that allows symptoms to be stored in a standard format. When problems occur, the Common Base Events in various WebSphere system logs (such as WebSphere Application Server and others) can be matched to known problems contained in the product's symptom database. Each product's symptom database can be downloaded from the IBM product support site from inside the Log and Trace Analyzer tool. In addition, you can create your own symptom databases to capture recurrent issues within your own IT organization.

Please click on the self-running demonstration inside this section to see these capabilities.

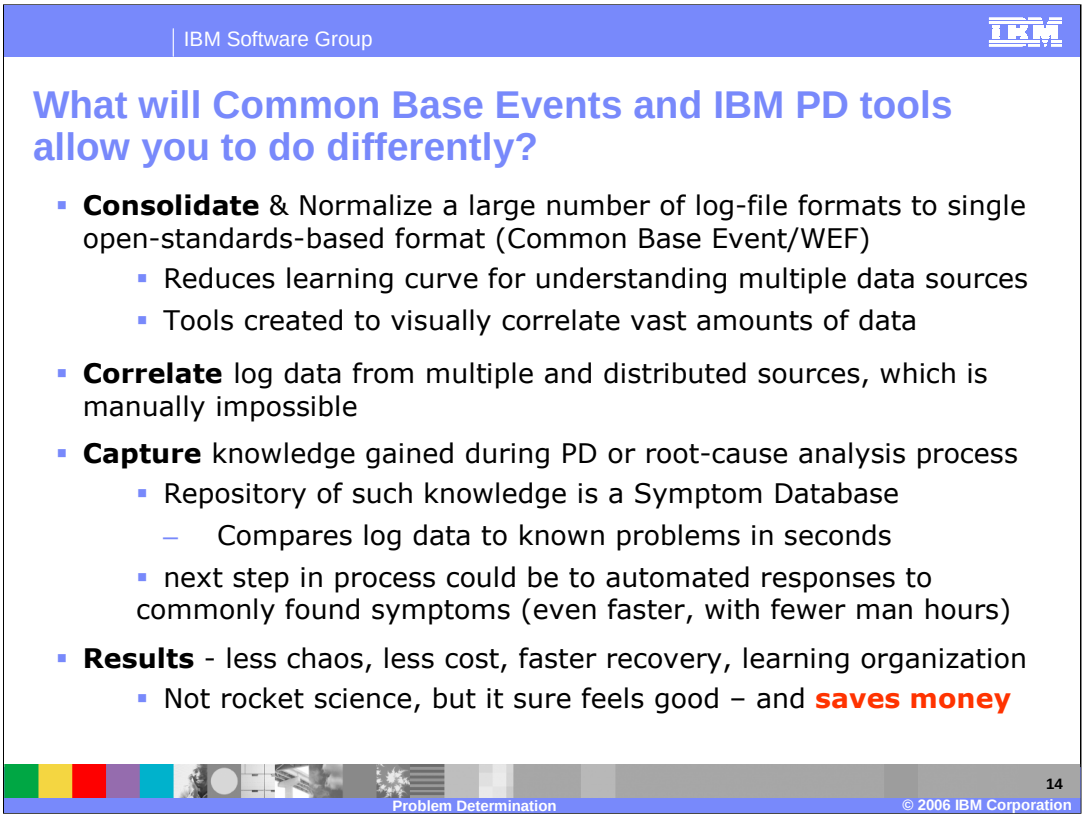

So, the common base event format and the problem determination tools just discussed allow you to do three very important things:

Consolidate a large number of different log files to a single standard format, which in turn allows you to ...

Correlate data from multiple different IT systems using tools that allow you to more quickly perform root-cause analysis, and then …

Capture the knowledge gained during the problem determination process by storing known problems in a symptom database, which later can be reused to solve known problems if they recur

The result is less chaos, finger-pointing, and time spent when your IT staff is performing problem determination, and faster times to system recovery and solving future issues. This has been proven in IBM customer engagements to save real money for organizations.

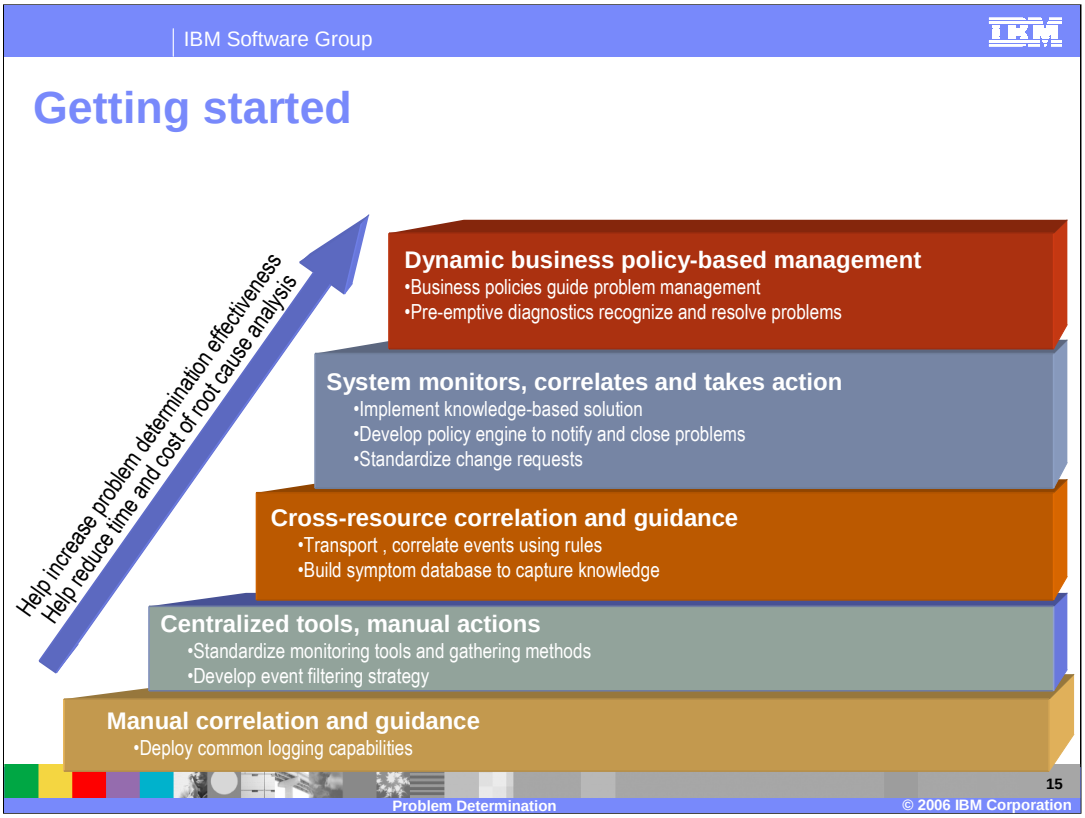

These problem determination capabilities within IBM software products can help your staff efficiently and effectively isolate the root cause of an infrastructure problem. But clearly, these capabilities represent only one piece of the puzzle.

Achieving best of breed problem management occurs through an evolutionary approach as shown on the screen. As you move through each layer of this roadmap, processes can be completed autonomically based on business policies so that your staff can focus on higher value tasks.

The first, basic level is manual analysis and problem solving. At this stage, it is critical to deploy common logging capability based on the WEF industry standard.

The second level supports centralized tools, but manual actions. This level requires standardized monitoring tools and data gathering process as well as the development of an event filtering strategy that identifies what information you need to isolate the source of a failure.

The third level supports cross-resource correlation and guidance. Here you can transport and correlate events from all components in your infrastructure, implement policies and rules for handling events, and begin capturing knowledge of past events and how they were resolved in a symptom database.

The fourth level represents implementation of a knowledge-based solution for monitoring and correlation of events across the enterprise, development of a rules policy engine that automatically notifies staff of problems and can close those problems with human intervention, and the standardization of the data models, and even the grammar, for change requests.

The fifth level represents a dynamic business policy-based management model. With this model, business policies guide problem management. Pre-emptive diagnostics automatically recognize and resolve problems. And call-home facilities are incorporated so that the latest fixes, patches and problem resolution information are downloaded into your symptom database automatically.

The topics discussed in this presentation are really how IBM is helping you take the first step in streamlining and simplifying some very basic, yet extremely tedious, problem determination processes. Embedding these technologies will not solve all the challenges you face in maintaining end-to-end service availability and in implementing all the stages displayed on this roadmap. However, the products and capabilities discussed will provide a vital piece of this puzzle.

Next steps

- Contact your local IBM representative for more information
- Visit **ibm.com/**autonomic/pd

If you are interested in more information about IBM Self-Managing Autonomic Problem Determination capabilities within IBM products or would like to learn more about how IBM can help you apply the problem determination roadmap, speak with your local IBM representative or visit our Web site at www.ibm.com/autonomic/pd. This Web site provides a comprehensive overview about our problem determination capabilities and which of our products currently incorporate these innovative technologies. Thank you for listening to this education module.

Template Revision: 04/25/2006 11:09 AM

Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM	CICS	IMS	MQSeries	Tivoli
IBM(logo)	Cloudscape	Informix	OS/390	WebSphere
e(logo)business	DB2	iSeries	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

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

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or program(s) described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.