Improving Client Stability with Proactive Problem Management

Refael Mizrahi
Project manager with Intel IT

Shachaf Levi
Support engineer with Intel IT

Jeff Kilford
Support engineering manager with Intel IT

Using our process, we reduced the number of blue screens by more than 3,000 per week.

Executive Overview

Intel IT implemented a proactive problem management process based on analysis of objective, largely system-generated data from client PCs across our worldwide environment. Using this approach, we have increased client stability by reducing the number of blue screen system crashes by more than 50 percent, and we are beginning to realize benefits in other areas including application crashes.

Traditionally, IT support has been reactive and emergency-driven; users report issues to the service desk, and the IT organization responds. However, this relies on subjective reports and focuses support only on the reported issues.

Our process differs significantly from this, adding a proactive approach:

- We regularly harvest data from the environment to identify existing and emerging issues. As part of this, we developed a tool that collects system crash dumps from thousands of clients and parses them to identify the root cause.

- We categorize and prioritize these issues based on impact, business value, and IT resources.

- We apply a problem management cycle to analyze the top-priority issues, deploy solutions, and measure the results.

Using our process, we reduced the number of blue screens by more than 3,000 per week. We have also begun to resolve application issues; this has improved security by reducing crashes of antivirus software.

Our data analysis enables us to identify users who are experiencing the most problems, even when they have not reported them; we proactively approach them and offer to fix their problems. Users have responded very positively to this. We are addressing other areas, including monitoring hard drives to detect problems before catastrophic failure; this helps us avoid the loss of data, increases employee productivity, and saves money because hard drives that are still under warranty are replaced at no cost to Intel.

Our experience suggests that proactive problem management—in tandem with deploying Intel® vPro™ Technology in our environment to improve manageability and reduce total cost of ownership (TCO)—can improve client stability, boosting users' productivity, overall satisfaction, and perception of IT, while reducing cost, enhancing security, and making more efficient use of IT resources.

Contents

IT@INTEL

IT@Intel is a resource that enables IT professionals, managers, and executives to engage with peers in the Intel IT organization—and with thousands of other industry IT leaders—so you can gain insights into the tools, methods, strategies, and best practices that are proving most successful in addressing today's tough IT challenges. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BACKGROUND

Intel's worldwide computing environment includes more than 100,000 enterprise PC clients. The environment is complex; it includes multiple client platforms and builds, and thousands of applications.

Maximizing client stability is extremely important. Users rely heavily on their PCs; system and application crashes have a substantial impact on user productivity and often result in the loss of unsaved work. These crashes also consume considerable IT support resources. In addition, if users experience frequent PC issues, they can develop a negative perception of an IT organization.

In late 2008, we were experiencing an average of about 5,500 "blue screen" system crashes in our client environment per week. As a result, users were not satisfied with the stability of their PCs.

We immediately began working to resolve this problem. At the same time, we realized that if we could reduce the incidence and impact of similar problems in the future, we could improve overall client stability and achieve long-term benefits as a result. Accordingly, we began developing a process aimed at achieving this more far-reaching goal.

This required a substantial change in our approach to client support. As at most IT organizations, Intel IT support traditionally has primarily been reactive and geared toward solving specific incidents. Users reported issues, and Intel IT fixed them.

Our challenge was to move from a reactive to a proactive mode. Rather than reacting only to user emergencies, we set out to develop a proactive problem management process that would enable us to identify major issues as they emerge within our environment, prioritize and allocate resources to these issues, and methodically address them before they impact all our users.

We realized that proactive problem management could deliver potential benefits, including greater user productivity, a better user experience resulting in an improved perception of IT, reduced IT support costs, and direct savings to Intel by identifying and replacing failing PC hardware components while they are still under warranty.

SOLUTION

Our approach was based on the Information Technology Infrastructure Library (ITIL), a set of concepts and policies that Intel IT has adopted for IT service management.

ITIL identifies the need for both reactive and proactive problem management. These play complementary roles within the enterprise.

Traditional Reactive Problem Management

This is the way that most IT organizations currently manage issues. It is driven by emergencies: Users experience events so destructive that they call the IT organization to help remediate the issue. The data recorded from these events is typically subjective, because it is based on input from users and support agents. Users' productivity is decreased until the issue is fixed. If several similar incidents occur and are recorded correctly, IT may be able to begin a problem investigation starting with the subjective data that has been collected.

One additional concern with relying entirely on reactive problem management is that users do not always report issues. Users may experience a system crash and a loss of productivity, and develop a negative perception of the IT organization as a result— yet the IT organization doesn't even know about it, because the organization is focusing only on the incidents that users do report.

Proactive Problem Management

Our vision of proactive problem management is an iterative process based largely on system-reported data. We harvest data from the environment at regular intervals and then sort the data so that we can identify priorities based on problems reported by the client systems. This data is far more objective than data recorded from service desk calls. In addition, because the data is gathered automatically, we are not dependent on users experiencing and reporting defects.

By gathering this data, we can gain an environment-wide view and analyze trends and emerging issues. This enables us to clearly identify the major issues that affect largest number of users—or soon will affect them. We can use this information to proactively address these major issues, rather than waiting for users to experience and report problems.

Goals

We set out several long- and short-term goals. These included:

- Reducing the number of problems that users experience. The initial target was to reduce the number of blue screens by 50 percent by the end of 2009.

- Implementing a more advanced, proactive problem management process across Intel's

IT organization, in keeping with our internal customers' expectations of a vigilant support organization. We aimed to be able to identify major issues before they affect large numbers of users.

- Reducing IT support cost through efficient use of IT support resources, by intercepting emerging issues rather than simply responding to emergencies.

- Our longer term vision was to monitor client health and improve overall stability—to monitor a range of potential issues so that we could predict and prevent problems. This meant focusing on more than system crashes; we also wanted to extend our approach to application crashes and hardware-related problems such as hard drives. We set a goal of reducing application crashes by 25 percent by the end of 2009.

Proactive Problem Management Process

We realized that we needed a broad process that we could use to analyze and resolve a range of issues. Defining and implementing this process required collaboration between multiple groups across Intel IT: operations, engineering, quality assurance, and release management.

Our process is shown in Figure 1. It consists of three main parts: initial collection of information to identify problems within the environment,

classification and prioritization of problems, and using a problem management cycle to resolve the problems identified as priorities.

COLLECT INFORMATION

Our process begins by collecting information to identify the problems that exist within the environment. To obtain the most complete picture of the issues across our environment, we collect information from various sources, including:

- The Intel IT incident management system
- Software agents installed on users' PCs
- User surveys and escalations

CLASSIFY AND PRIORITIZE THE PROBLEM

We log, categorize, and prioritize each problem in order to determine where to focus IT resources.

- **Log.** We log each new problem and add it to the list of problems in the pipeline.

- **Categorize.** We place each problem into the appropriate category, such as hardware or OS. The approach to resolution may differ depending on the category.

- **Prioritize.** To decide where to invest our resources, we analyze the problem impact, the resources needed to resolve it, and the business value delivered by solving the problem. The outcome is a list of the priority problems that we will work on.
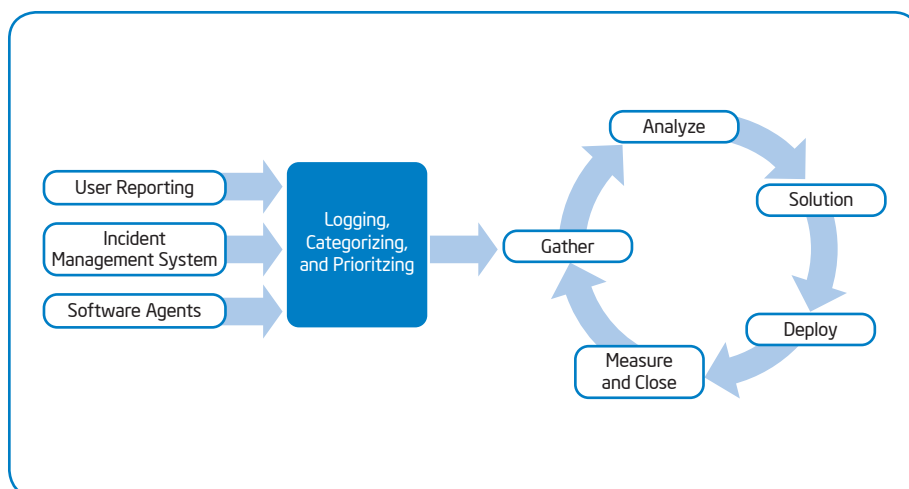

Figure 1. Proactive problem management process.

Intel IT Blue Screen Dump Collector Tool

When a PC experiences a blue screen system crash, a dump file is created on the system. However, because it resides on the user's PC, the information in the dump file is not automatically available to the IT organization. Unless the user contacts the support desk, the IT organization may not know about the problem.

We developed a tool that collects the dump files from each PC and parses them to identify possible root causes and other key information. Then the information is uploaded to a central database for analysis.

This enables us to gather more accurate and objective data. For example, we can measure the number of crashes that actually occurred, rather than just the ones that were reported. We can also analyze the true cause of the crashes, rather than relying on information that was recorded during user calls to the service desk.

MANAGE THE PROBLEM

For each of these priority problems, we use a five-step problem management cycle.

1. **Gather.** The team works to gather all the information relevant to the problem so that we can analyze it. The data may reside on thousands of client PCs, requiring tools to collect it from the environment. Because we could not find an existing tool to do this for blue screens, we developed one ourselves (see sidebar). All the relevant data is uploaded into a central database.

2. **Analyze.** We analyze the data and find the root cause for the problem, using tools including a client environment dashboard we created. This analysis enables us to answer questions such as: Which driver causes the most blue screens? On which PC models is the problem occurring? Which driver version is installed on those PCs? Is a newer version available? Is the problem restricted to a specific building location?

3. **Find a solution.** Once we have identified the root cause of the problem, we look for a permanent solution—and a temporary workaround if necessary. For example, if the root cause is an outdated version of a driver, we check to see whether the new version will solve it. If, on the other hand, the root cause is the current version of a driver, we may need to escalate the issue to the supplier for resolution.

4. **Deploy the solution.** Once we have the solution, we need to decide how to deploy it into the production environment. We have several choices. We can deploy the fix to:

   – Only those users who experienced the issue.

   – All clients. This is the most proactive approach because it can fix the issue before most users experience it. However, there is a risk that the fix can create additional problems.

   – A specific notebook model.

5. **Measure and close.** We follow up to determine the actual effect of our actions on the environment. We have several ways to do this: We can survey users and agents, check the incident management tools, and collect data directly from the clients.

GAP ANALYSIS

We initially performed a gap analysis to determine which support areas we need to improve in order to implement the process. We examined several areas, including the infrastructure required to gather, store, and report issues; data analysis tools; solution deployment capabilities; and process governance, including the indicators that verify progress.

We identified several gaps. For example, we lacked a process to automatically deploy a fix to one specific group of users, and we worked with the group to create a process that can be used by our team and by others.

We also realized that we needed better reporting and business intelligence capabilities in order to analyze the data. To fill this need, we created our client dashboard. This displays on a single page the key indicators we need, such as the total number of blue screens, the top causes, and the trend over the past year. In the longer term, we plan to create a data warehouse and business intelligence portal to enable more detailed analysis.

Results

We first applied our process in late 2008, focusing on resolving system crashes after the number of blue screens surged to 5,500 per week. Since then, we have begun to address a broader set of issues.

BLUE SCREENS

Our initial target was to reduce the number of blue screens by 50 percent by end of 2009, from the baseline of 5,500 blue screens per week.

We gathered and analyzed information from clients using our dump collector tool. Our client dashboard showed that one driver was responsible for a significant proportion of the problem. Blue screens caused by this driver had risen from an average of 4 percent to 10 percent of the total.

After analyzing the data, our engineering teams and a software supplier provided a solution. We decided to deploy the fix to the affected users only. A week later, we were able to measure and close the issue by once again gathering data from PCs: Blue screens caused by the driver were back to the normal level.

Following this success, we applied several more problem management cycles to address other top root causes, shown in Figure 2, including problems related to network cards, remote connectivity, and display drivers.

By the end of the first quarter of 2009, we had reduced the number of blue screens from 5,500 per week to about 3,500—a reduction of approximately 35 percent. As shown in Figure 3, we subsequently were able to continue reducing the weekly total, reaching a 55 percent reduction during Q3 compared with the original 5,500 per week baseline. This beat our original goal of a 50 percent reduction by the end of 2009. We are now preventing about 3,000 client crashes weekly, resulting in an improved experience, greater productivity, and reduced data loss for our internal customers.

APPLICATION CRASHES

We have begun to apply our problem management process to reducing application crashes.

We use a tool from a software supplier to gather information about application crashes; the tool enables clients to send error reports to a central location. Based on an analysis of this data, we determined that about 85,000 crashes occur in an average week. About 65 percent of these crashes occur with 10 common applications, including a Web browser, antivirus software, an e-mail client, and search indexing software.

We have already seen some significant benefits. For example, we detected that our virus detection software generated a huge spike of about 70,000 additional crashes. We discovered the problem through our data analysis, even though users were not yet widely reporting it to the service desk. We formed a team dedicated to the problem, analyzed it, identified a fix, and deployed it, eradicating the problem in about two weeks with almost no disruption to the operations of our internal customers. We have subsequently deployed another fix designed to reduce other crashes with the same software.

As a result, by the third quarter of 2009 the number of application crashes had decreased by an average of about 15 percent.

A key advantage of our approach is that we have detailed, objective data that is almost


Figure 2. By collecting system information directly from users' PCs, we are able to monitor trends in the top causes of blue screens.

Number of Blue Screens Per Week

Figure 3. With proactive problem management, we reduced the weekly number of blue screens by about 55 percent by Q3 2009.

[1] The number of blue screens dipped temporarily at the end of 2008 due to employee vacations.

irrefutable. This enables us to present a very powerful case to software suppliers, which has resulted in increased efforts to analyze and fix software products.

This also helped us build a case internally to keep all products up to date, since in many cases suppliers will not work to resolve issues if most of our users have an outdated version of their product.

Another benefit is improving the quality of internally developed software. We saw that an internally developed application started to crash, and have reported the problem to the developer together with data that can be used to fix the issue and improve code quality.

HARD DISK DRIVES

Our enterprise PCs run health monitoring software, including a tool that checks for disk errors. This enables us to monitor the health of hard disk drives (HDDs), with the potential to replace them before they fail. Our IT engineering group has defined a disk error threshold; if a disk is experiencing more errors than this threshold, it is considered at risk of failure.

We have begun proactively contacting users whose PCs are experiencing large numbers of disk errors and offering to replace their hard drives. This helps us avoid catastrophic disk crashes that result in serious loss of data and user productivity. In addition, this has a direct financial benefit, since some of these hard drives are still under warranty and are replaced at no cost to Intel.

PROACTIVELY WORKING WITH USERS

Our data analysis enables us to identify the users who experience the most crashes. We have begun using this information to proactively approach these users and offer to work to resolve their issues, including a new build if needed.

To date, the feedback has been extremely positive. Our internal customers love the fact that we know the problems they are experiencing—even when they haven't reported them—and are offering to resolve the issues. This helps create a very favorable view of the IT organization.

FUTURE PLANS

We plan to continue to expand our approach, potentially addressing new aspects of the client environment and using new tools to solve problems.

Intel® vPro™ Technology

Our proactive management process is an element of our broader efforts to improve manageability and reduce total cost of ownership (TCO). As part of this, Intel IT is currently deploying PCs with Intel vPro technology throughout our environment. Intel vPro technology includes out-of-band capabilities that improve our ability to remotely manage clients, even when they are powered off or their OS is unresponsive. To date, we have focused on several specific use cases, including remote configuration and remote diagnosis and repair.

In the future, Intel vPro technology could potentially be used to support our proactive

problem management process, performing functions such as remotely rebuilding a system or replacing and upgrading drivers. Additional Intel vPro technology capabilities such as Fast Call for Help and Agent Presence Checking could help us extend our reach beyond the corporate firewall for diagnosis and repair of mobile PCs—and help to avoid issues from occurring.

Monitoring Batteries

We plan to monitor notebook batteries, using an approach similar to the way that we already monitor hard drives. We can check if a user's battery has lost a significant amount of its capacity and contact the user to replace the battery. This may result in a direct financial benefit to Intel: If the battery is still under warranty, it will be replaced at no charge.

Identifying and Refreshing Problem-prone PCs

We are implementing business intelligence tools that will enable us to use the collected information to better model problems and take actions as a result. For example, we could identify PCs that are particularly prone to problems and target them for early refresh.

We can do this by establishing a problem threshold, which could be based on a combination of problems that we monitor, such as application crashes and blue screens. PCs that cross this threshold could be refreshed before the end of the standard refresh cycle.

This approach could lead to increased productivity and reduced TCO as support costs decline.

CONCLUSION

Proactive problem management is already delivering several benefits.

These include:

- **Greater system stability.** This leads to increased user satisfaction.

- **Increased user productivity.** We have reduced the number of blue screens by more than 3,000 per week; this means more than 3,000 fewer reboots requiring several minutes each. This adds up to a sizable productivity gain across the enterprise. In addition, there is an associated reduction in the amount of data and work lost due to crashes.

- **Positive customer feedback.** When we proactively approach users who have been experiencing issues and offer to fix their problems, the feedback is extremely positive.

- **Improved security.** When antivirus or other system-protection software crashes, the client is left unprotected, providing a potential enterprise entry point for malware. Adding to this concern is the fact that these application crashes may have been caused by malware, which often attempts to attack the antivirus software on a system that it infects. When we reduce the incidence of antivirus application crashes, the result is increased enterprise security.

- **Software quality.** The information we collect is being used to improve the quality of internal as well as external products, helping to make IT a strategic partner to product development groups.

We expect to realize additional benefits as we continue to use and expand the process.

ACRONYMS

HDD	hard disk drive
ITIL	Information Technology Infrastructure Library
TCO	total cost of ownership

For more straight talk on current topics from Intel's IT leaders, visit www.intel.com/it