



*The knowledge
behind the network.®*

A Framework for Incident and Problem Management

By Victor Kapella

Consulting Manager

International Network Services



A Framework for Incident and Problem Management

By Victor Kapella, Consulting Manager

Introduction

Many organizations have developed multi-tiered, information technology (IT) support services delivered by help desks, network operations centers (NOCs) and engineering organizations. A common mistake made when developing these services is to focus on responding to incidents instead of on preventing problems from occurring in the first place. The relationship among these service activities is not well understood, thus many organizations fail to successfully execute proactive problem prevention.

This whitepaper defines incident and problem management based on the Information Technology Infrastructure Library (ITIL) Service Support best practices and INS's experience in the industry. It further explains the differences between incident management and problem management and offers a framework for addressing both activities.

The Language of Incident, Problems and Errors

The ITIL Service Support is an internationally recognized best practices model used to guide IT organizations in developing their service management approaches. This model has been widely adopted. It is prescriptive in nature and identifies elements, in addition to incident and problem management, that need to be addressed to successfully run an IT organization like a service business. This model defines a technical vocabulary for the discussion of support services. It defines clear concepts and draws distinctions between various support activities. For example, the activities required to respond to service interruptions and to restore service have different qualities than those activities required to identify and permanently remove the underlying cause of service interruptions.

Incidents

An incident is any event that is not part of the standard operation of a service and causes, or may cause, an interruption to or reduction in the quality of that service. Examples of incidents are:

- ▶ User cannot receive e-mail
- ▶ NOC monitoring tool indicates that a WAN circuit may be down
- ▶ User perceives that an application is running slow

Problems

A problem is an unknown, underlying cause of one or more incidents. A single problem may generate several incidents.

Errors

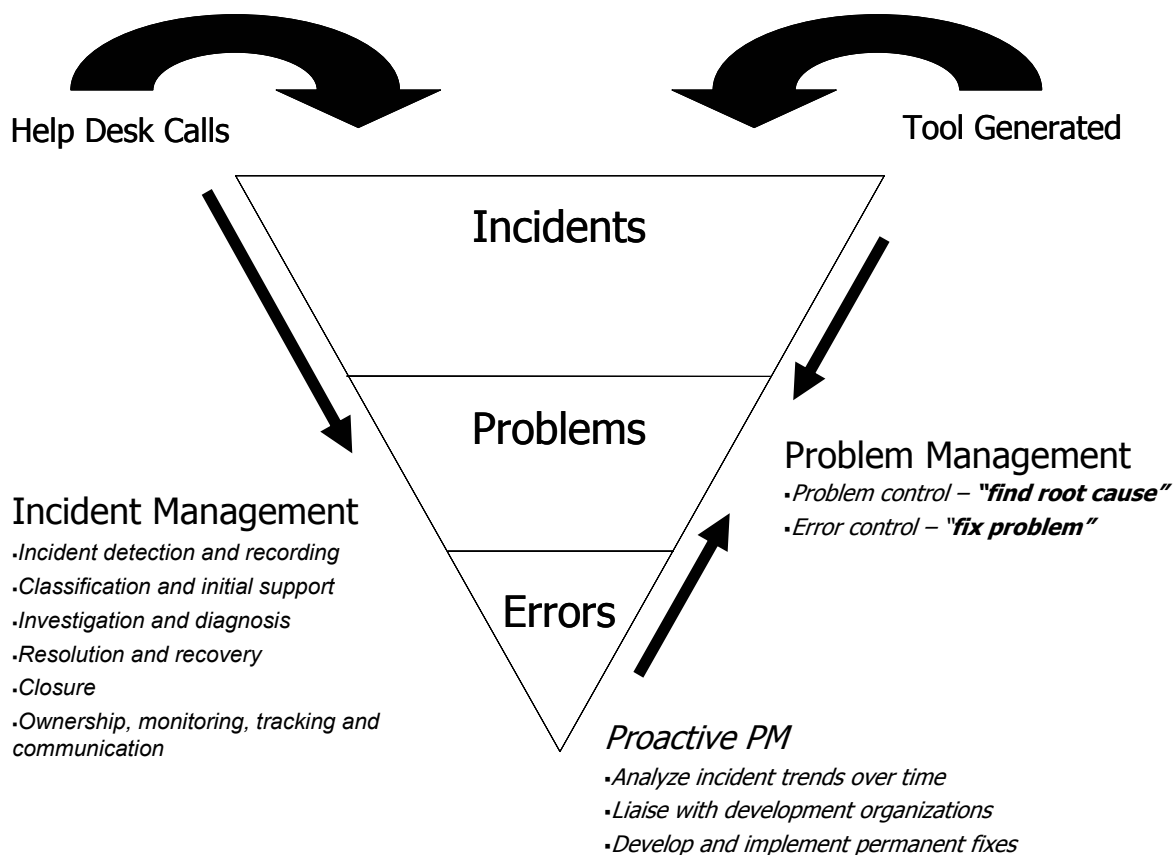
An error is a problem for which the root cause has been identified and a workaround or permanent solution has been developed. Errors can be identified through analysis of user complaints or by vendors and development staff prior to production implementation. Examples of errors include:

- ▶ Laptop network settings misconfigured
- ▶ Monitoring tool misidentifies WAN circuit status when polled router is busy

Managing Incidents and Problems

The key concepts and language of incident and problem management are shown in Figure 1. There is a lifecycle relationship among incidents, problems and errors: incidents are often the indicators of problems; problems lead to the identification of the root cause of the underlying error; errors are then systematically eliminated.

Figure 1: IM and PM Concepts



Incident Management

Incident management (IM) refers to activities undertaken to restore normal service operation as quickly as possible while minimizing adverse impact on business operations. IM is a reactive, short-term focus on restoring service. IM activities include:

- ▶ Incident detection and recording
- ▶ Classification and initial support
- ▶ Investigation and diagnosis
- ▶ Resolution and recovery
- ▶ Closure

Problem Management

Problem management (PM) refers to activities undertaken to minimize the adverse impact on the business of problems that are caused by errors within the IT infrastructure, and to prevent recurrence of incidents related to these errors. PM gets to the root cause of problems, identifies workarounds or permanent fixes and eliminates errors. PM activities include:

- ▶ Problem control
- ▶ Error control
- ▶ Proactive problem prevention
- ▶ Major problem reviews

Problem Control

The purpose of problem control is to find the root cause of a problem by executing the following steps:

- ▶ Identifying and recording of the problem
- ▶ Classifying the problem and prioritizing response activities
- ▶ Investigating and diagnosing root causes

Error Control

Error control activities ensure that problems are fixed by executing the following steps:

- ▶ Identifying and recording known errors
- ▶ Assessing permanent fixes and prioritization
- ▶ Resolution recording of temporary workarounds into service support tools
- ▶ Closure of known errors by implementing permanent fixes
- ▶ Monitoring known errors to determine if a change in priority is warranted

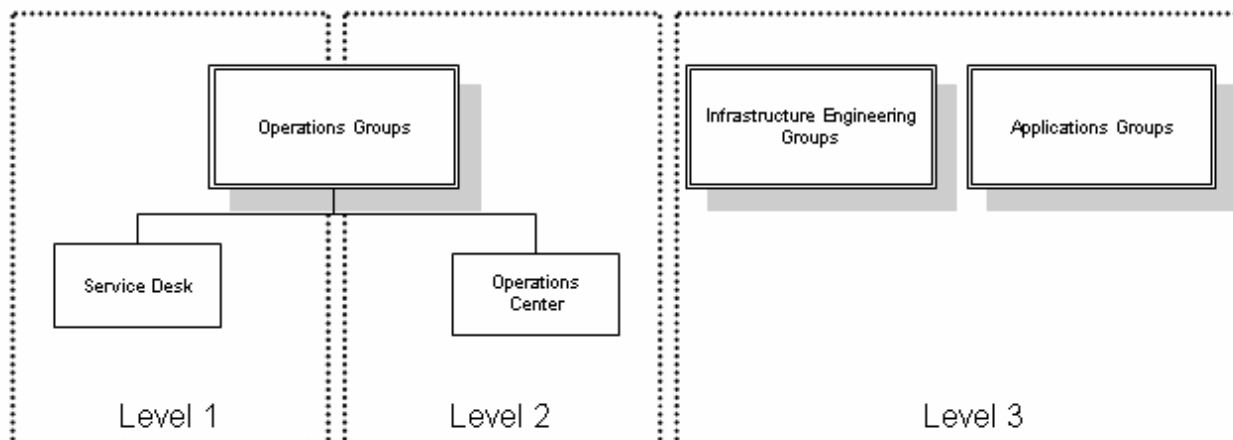
Problem Review

The purpose of a problem review is to improve IM and PM processes. This is accomplished by performing a post-mortem examination of the quality of the IM and PM response activities associated with a major incident or problem.

Organizational Roles and Responsibilities

The most common support structure that INS encounters is a tiered model where increasing levels of technical capability are applied to the resolution of an incident or problem. A typical organizational structure for this tiered support model is shown in Figure 2.

Figure 2: Typical Tiered Support Model



The actual roles and responsibilities seen in tiered support implementations are as varied as the people, history and politics that comprise an organization's environment. The following description of a tiered support model is typical in many organizations.

Level 1 Support

The organization providing level 1 support commonly resides in the Operations group and is typically identified as a Call Center, Help Desk, Service Desk or other similar name.

Roles

Owner of the IM process. Level 1 support ensures that a well defined, consistently executed, properly measured and effective IM process is established and maintained.

Receive and manage all customer service issues. Level 1 support is the single point of contact for reporting service issues, and acts as end-user advocate to ensure that service issues are resolved in a timely fashion.

First line of support. The level 1 organization makes the first attempt to resolve the service issue reported by the end user.

Responsibilities

Accurately record incidents. Level 1 support ensures that an incident is properly logged into the incident management system. In doing so, it must:

- ▶ Ensure that the ticket contains an accurate and properly detailed description of the problem
- ▶ Ensure that the severity/priority classification is correct
- ▶ Determine the nature of the problem, business partner contacts, impacts and expectations

Own every incident. As the end-user advocate, level 1 support owns the successful resolution of every incident. It ensures that the IM process resolves the issue in a timely fashion by:

- ▶ Developing and managing a resolution action plan
- ▶ Initiating specific assignments for staff and business partners
- ▶ Escalating the incident as required when resolution targets are missed
- ▶ Ensuring internal communication occurs according to defined service targets
- ▶ Championing the interests of the involved business partners

Level 1 support uses the problem management database to match incidents with known errors and to apply previously identified workarounds to resolve incidents. Its target is to resolve 80 percent of incidents. The remaining incidents are escalated to level 2

Continually improve the IM process. As owner of the IM process, level 1 support ensures that the process and capabilities are adequate, and are improved when necessary by:

Evaluating the effectiveness of the IM process and supporting mechanisms such as reports, communication formats/messages, and escalation procedures

- ▶ Developing department-specific reports and procedures
- ▶ Maintaining and improving communication and escalation lists
- ▶ Participating in the problem review process

Capabilities

Interpersonal skills paramount; technical skills secondary. Level 1 support personnel are primarily involved in triage and management of problems. Very little technical troubleshooting should occur at this level of support.

Ability to apply “canned” resolutions. Level 1 personnel should have the ability to recognize patterns of symptoms, apply search tools to identify previously developed solutions, and help end-users implement the solution.

Level 2 Support

Also typically residing in the Operations group, level 2 support organizations are commonly called Command Centers, Network Operations Centers, or Distributed Computing Control Centers.

Roles

Troubleshoot incidents. Level 2 support investigates, diagnoses and resolves most incidents that are not cleared by level 1 support. These incidents tend to be indicative of new problems.

Owner of PM process. Level 2 support ensures that a well-defined and effective problem management process, as previously defined, is in place.

Proactive management of the infrastructure. Level 2 support uses tools and processes to ensure that problems are identified and resolved before incidents occur.

Responsibilities

Resolve incidents escalated from level 1. Whereas level 1 is expected to resolve 80 percent of incidents, level 2 support is expected to resolve 75 percent of incidents that are escalated to them, for an overall total of 15 percent of the incidents reported to level 1 support. The unresolved incidents are escalated to level 3 support

Determine root cause of problems. Level 2 support determines the root cause of problems and identifies workarounds or permanent fixes. They engage and manage other resources as necessary to determine the

root cause. They escalate problem resolution to level 3 support when the root cause is an architectural or technical issue that exceeds their skill-set.

Champions the implementation of workarounds and permanent fixes. Level 2 support ensures that projects are raised within development organizations to implement permanent fixes to known errors. They ensure that workarounds are documented and communicated to level 1 support staff and implemented in tools.

Proactively monitor the infrastructure. Level 2 attempts to identify problems before incidents occur by monitoring infrastructure components and taking corrective action when defects or problematic trends are discovered.

Proactively examine incident trends. Past incidents are examined to determine if there are underlying problems that need to be fixed before future incidents occur. Incidents that are closed without being matched to a known problem are also examined for potential underlying problems.

Continually improve the PM process. As owner of the PM process, level 2 support ensures that the process and capabilities are adequate, and improved when necessary. They lead the problem review process to determine lessons learned and ensure that process controls, such as meetings and reports, are adequate.

Capabilities

Technically competent with reasonable interpersonal skills. Level 2 support staff should have a range of technical skills across the technologies that are supported, including networks, servers and applications. A common deficiency in level 2 support organizations is in operating system or application expertise. Further, there should not be a significant skill gap between the level 2 and level 3 organizations. Some level 2 staff should be as skilled as level 3 support staff.

Network, server and application knowledge. The level 2 organization needs to be able to resolve incidents and problems across the gamut of technologies at use in the organization.

Level 3 Support

This level of support typically resides in the Engineering and/or Development groups within IT. These organizations are commonly called Engineering, Architecture, Network Integration, or Applications Development.

Roles

Planning and design of IT infrastructure. Typically, the level 3 support group has a minor role in IM and PM as these organizations are chartered primarily with planning and design of the IT infrastructure. As such, their goal is to implement defect-free infrastructure that is not the source of problems and incidents.

Last escalation group. If an incident or problem exceeds the technical capabilities of the level 2 support group, the level 3 support group takes responsibility to reach resolution.

Responsibilities

Resolve incidents escalated from level 2. As most incidents are caused by previously known errors, very few incidents (5 percent) should percolate up through level 2 support and into the level 3 organization. Level 3 is responsible for resolving all incidents that percolate up.

Participate in PM activities. Level 3 support are involved in finding root cause, workarounds, and permanent fixes.

Implement permanent fixes to remove errors from the infrastructure. Level 3 has a significant role in planning, designing and implementing projects that provide permanent fixes to the infrastructure. These projects must be prioritized along with the normal development work to achieve a desired balance.

Capabilities

Subject matter experts. Level 3 support teams should be subject matter experts who plan and design the IT infrastructure.

Processes

There are three fundamental processes in the framework: Incident Management, Problem Control and Error Control. These basic processes exist in best-in-class organizations. They may not be called specifically by these names, but they are always present.

Incident Management Process

The IM process focuses on restoring interrupted service as soon as possible. Table 1 describes the IM process elements. Figure 3 shows a diagram of how the IM process works and its linkage to PM.

Table 1: Incident Management Process Elements

Process Element	Description
Purpose	Restore service to end user while maintaining high satisfaction
Owner	Level 1 support team
Input	User call to report service interruption
Output	Service restored End user notified Incident record created Possible problem record created
Typical Measurements	Quantity of tickets presently open by severity, longest elapsed time, responsibility group Quantity of incidents by time (monthly/quarterly) Quantity of tickets escalated and resolved by each support level Mean elapsed time tickets were assigned in each group Mean time to restore service Percentage of incidents resolved by resolution timeframe target Tickets generated by technology Tickets generated by user group

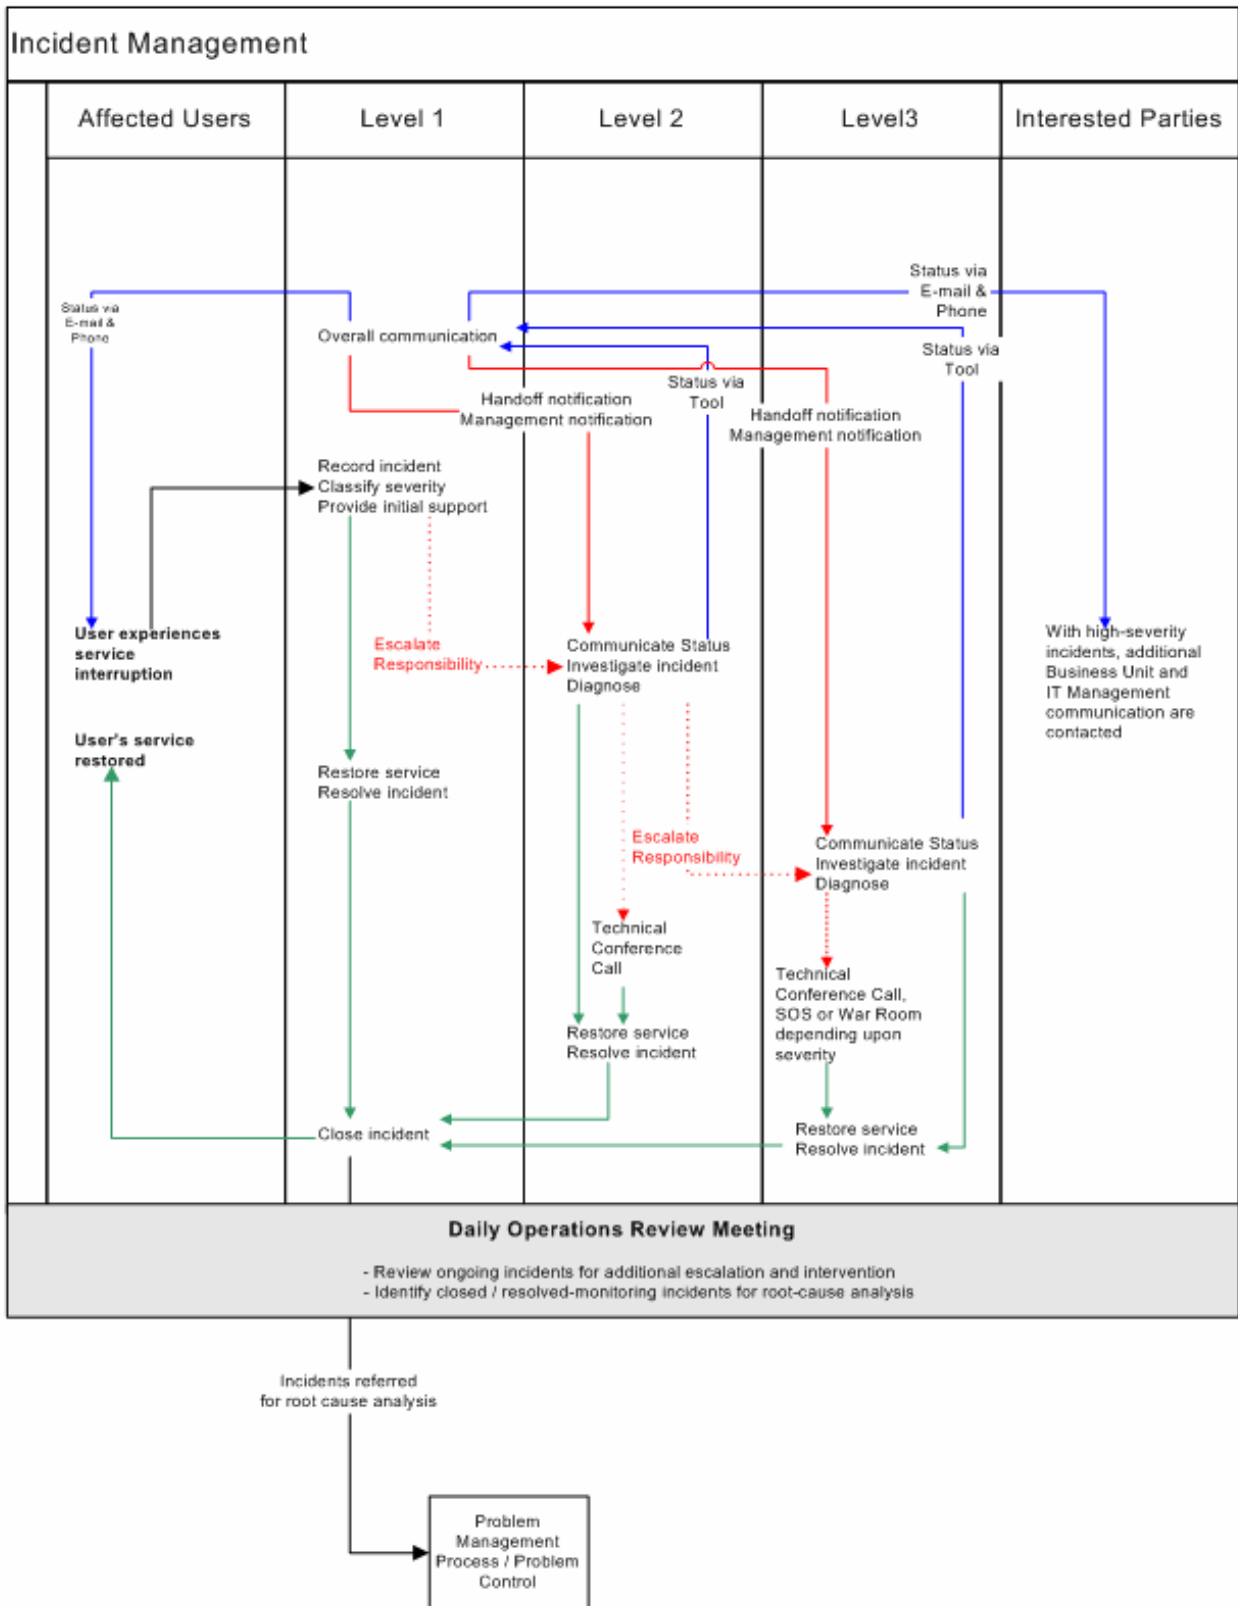
This process model illustrates the principals of IM as defined by ITIL and through observation. Several process activities are worth mentioning and will be discussed in greater detail in later sections of this document.

First, all communication is well defined. That is to say, there is a clear distinction between communication intended to convey status and communication intended to request action. Communication occurs at a minimum according to a set schedule. The people and organizations requiring communication are defined in advance of the incident. Further, the content of action requests and status communications are standardized.

Second, escalation is well defined and executed according to a set schedule. Responsibility for incident resolution is progressively passed to more capable organizations until service is restored.

Third, process control is maintained through a daily “top 10” meeting, where open and exceptional incidents are reviewed for further escalation or referral to PM activities.

Figure 3: Incident Management Process Model



Problem Control Process

The Problem Control process focuses on prioritizing, allocating and monitoring efforts to determine the root cause of problems and to identify workarounds or permanent fixes. This process could be likened to program management, where each problem is a project that has to be managed along with a portfolio of other such projects. The basic elements of the problem control process are shown in Table 2.

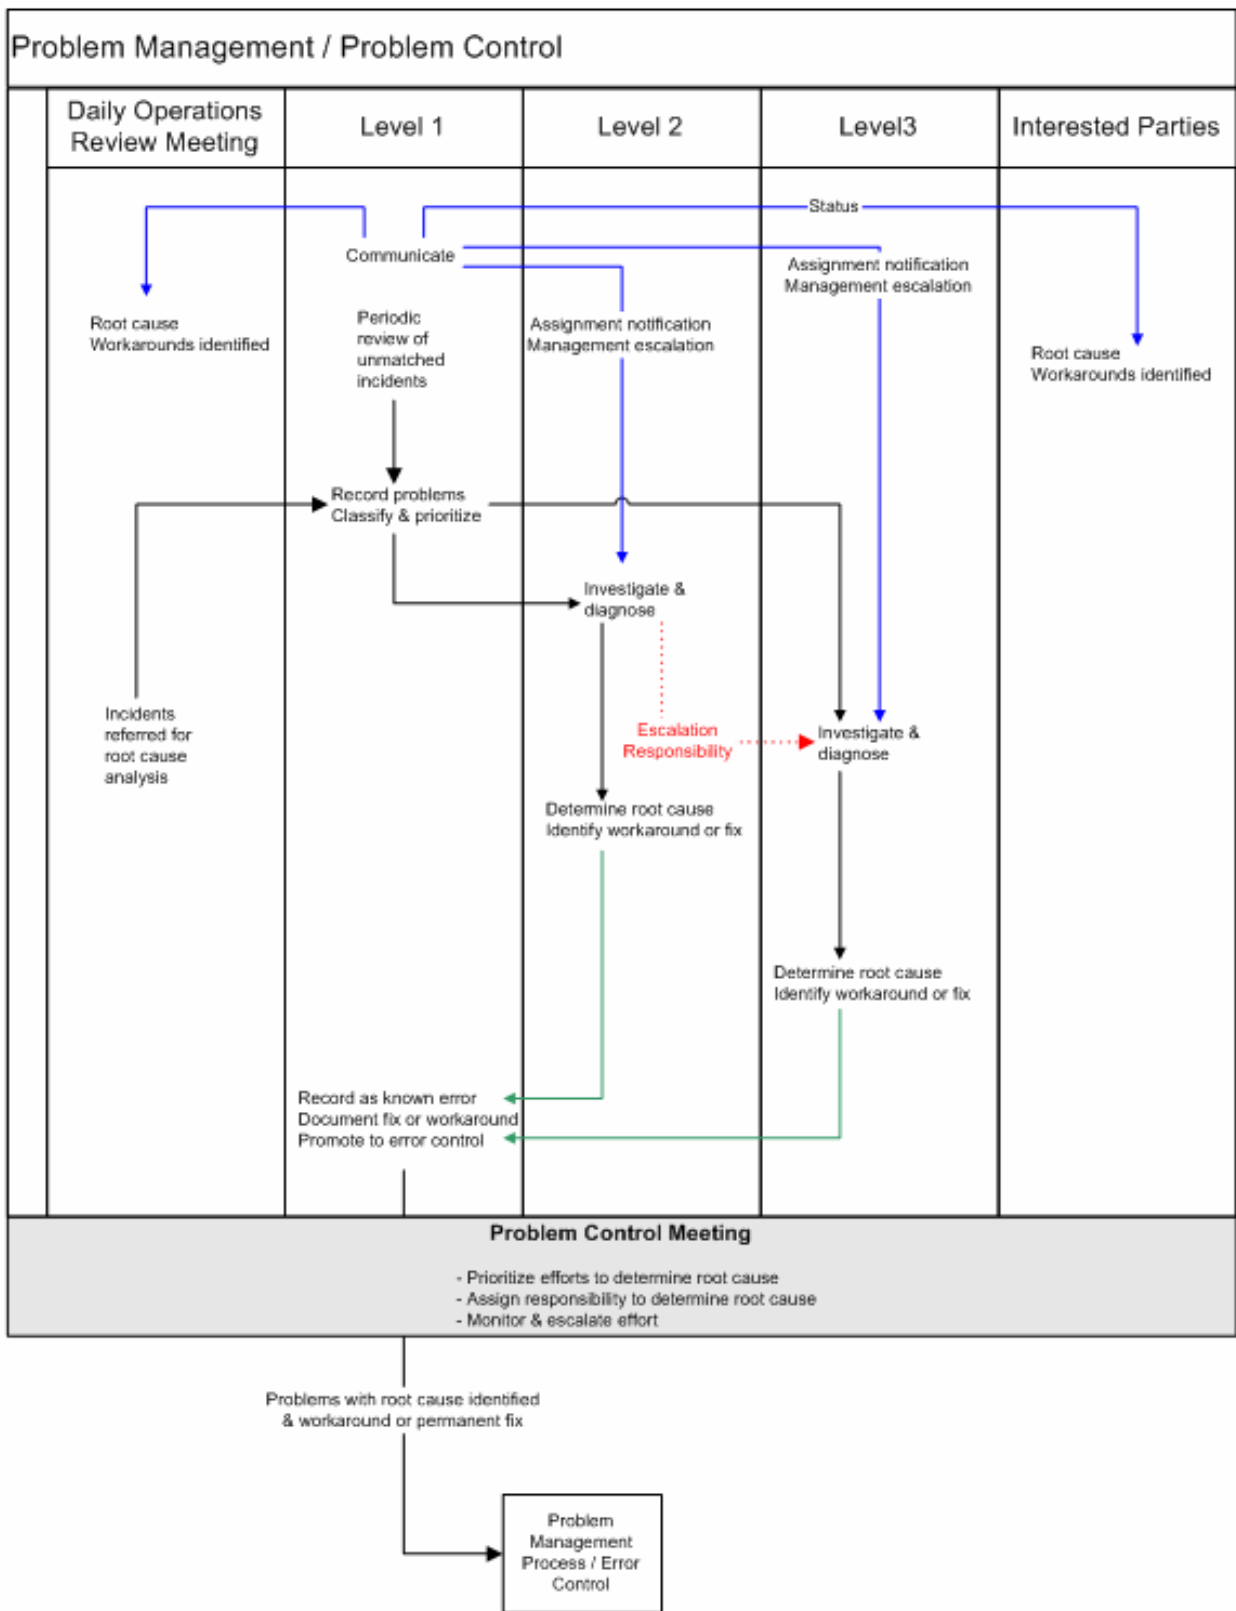
Table 2: Problem Control Process Elements

Process Element	Description
Purpose	Determine root cause of problems and identify temporary workarounds or permanent fixes
Owner	Level 2 support team
Input	High-severity incidents Incidents referred to level 3 support for resolution Incidents referred by the “top 10” meeting
Output	Documented root cause Communicate temporary workarounds to all support levels
Measurements	Quantity of problem tickets referred over time (monthly/quarterly) Quantity of problems where root-cause analysis is waived Quantity of problem tickets presently open (without root cause identified) Quantity of problems assigned to level 2 & level 3 Mean time tickets were assigned to each level Mean time to determine root cause Tickets generated by technology Tickets generated by user group

Input into Problem Control can come from several sources. Typically, the highest severity incidents are automatically submitted to the problem control process. In organizations that have robust level 2 support, incidents that are referred to level 3 are routinely forwarded to problem control also. And finally, the daily operations review meeting refers incidents to problem control.

The process for implementing problem control is shown in Figure 4.

Figure 4: Problem Control Process Model



The focus of the Problem Control process is the identification of the root cause. Participants in a root-cause analysis and the length of time to complete such an analysis are as varied as the problems themselves. It is fair to state the following:

- ▶ If you have a sufficient quantity of problems, appoint a standing team. Otherwise, raise a team as problems are referred for root-cause analysis much like you would form a team to address any project.
- ▶ The team will almost certainly require cross-disciplinary expertise. But this is dependent upon the nature of the problem being referred.
- ▶ The length of time to determine the root cause should be estimated (project plan developed) when the problem is referred to the team. The team's progress should be measured against this estimate.

Once resources are allocated and prioritized, the actual mechanics of determining the root cause can take several forms. Examples of best-practice, root-cause techniques include Kepner and Tregoe Analysis and Ishikawa Diagrams.

Kepner and Tregoe Analysis

The method developed by Charles Kepner and Benjamin Tregoe to analyze problems states that problem analysis should be a systematic process of problem solving and should take maximum advantage of knowledge and experience. According to their method, there are five phases for Problem analysis.

1. Define the Problem
2. Describe the Problem with regard to identity, location, time and size
3. Establish possible causes
4. Test the most probable cause
5. Verify the true cause.

These phases are relevant to all root-cause analysis efforts. The degree of effort and thoroughness of each phase will vary depending upon prioritization and available information.

Low severity problems may experience a corresponding low level of effort, skill level and priority for determining root-cause. Whereas a high-severity problem may merit appointment of a cross-disciplinary team, much like the investigation of an air-crash, with high-level of priority, effort and established time-frames for determining root cause.

In either event, the methodology is the same. It is the structured approach that helps ensure success in determining the root-cause.

Ishikawa Diagrams

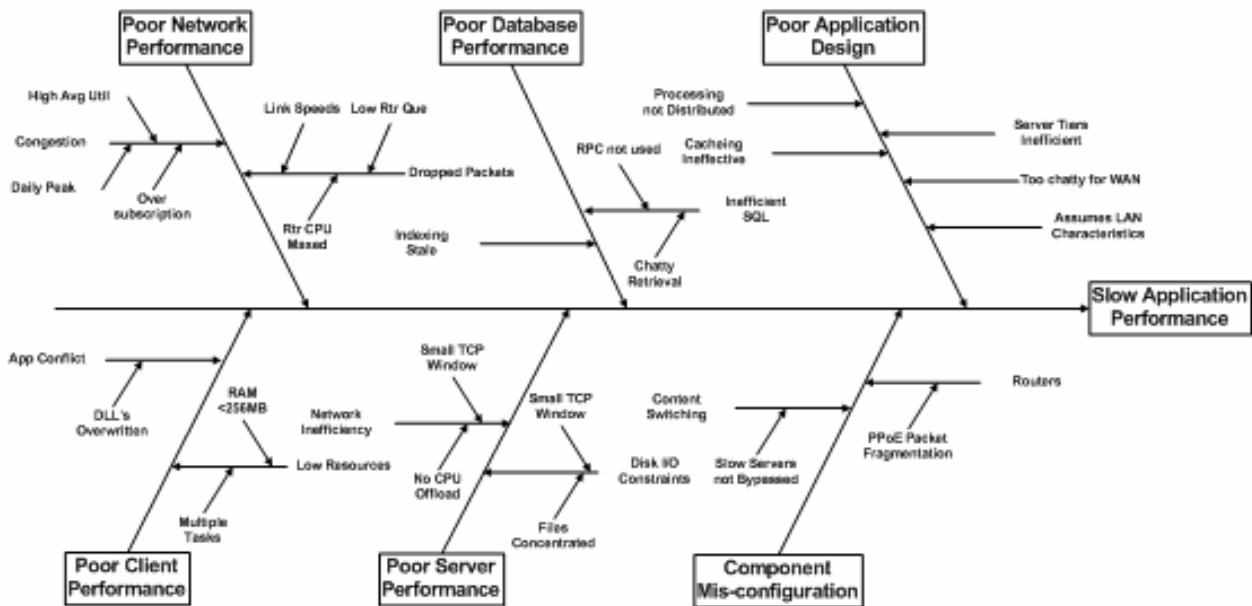
The Ishikawa diagram, also referred to as a cause-and-effect diagram, attempts to identify the root cause of a problem by diagramming a series of causes and effects that relate to the problem that was experienced. The Ishikawa diagram can achieve the following:

- ▶ Focus attention on one specific problem
- ▶ Organize and display graphically the various theories about the root causes of a problem
- ▶ Show the relationship of various factors influencing the problem
- ▶ Provide an aid for problem solving.
- ▶ Reveal important relationships among various variables and possible causes
- ▶ Provide additional insight into process behaviors

- ▶ Focus the team on the causes, not the symptoms

Ishikawa diagrams are used to facilitate a brainstorming session of a group of experts during which members identify primary and secondary factors which may have led to the problem. The main problem, as shown in Figure 5, is represented by the trunk of the diagram, and primary contributing factors are represented as branches. Secondary factors are then added as stems, and so on. Creating the diagram stimulates discussion and often leads to increased understanding of a complex problem.

Figure 5: Example Ishikawa Diagram




Prioritizing Problems

Reports, such as quantity of incidents by technology, can be used to develop Pareto charts, which are useful tools for prioritizing root-cause efforts. Pareto charts help visualize the Pareto principle, which states that a small subset of problems tend to occur much more frequently than the remainder. A common expression of the Pareto principle is “the 80/20 rule.” Addressing the important 80 percent of an issue is more effective than trying to address 100 percent.

A Pareto chart can be used to decide which subset of problems should be solved first, or which problems deserve the most attention. Pareto charts can also be used to provide a before-and-after comparison of the effect of control or quality improvement measures. The Pareto chart in Figure 6 shows that the root cause of approximately 80 percent of incidents reported to the helpdesk are due to non-standard desktop configurations and non-specific application failures.

Figure 6: Example Pareto Chart



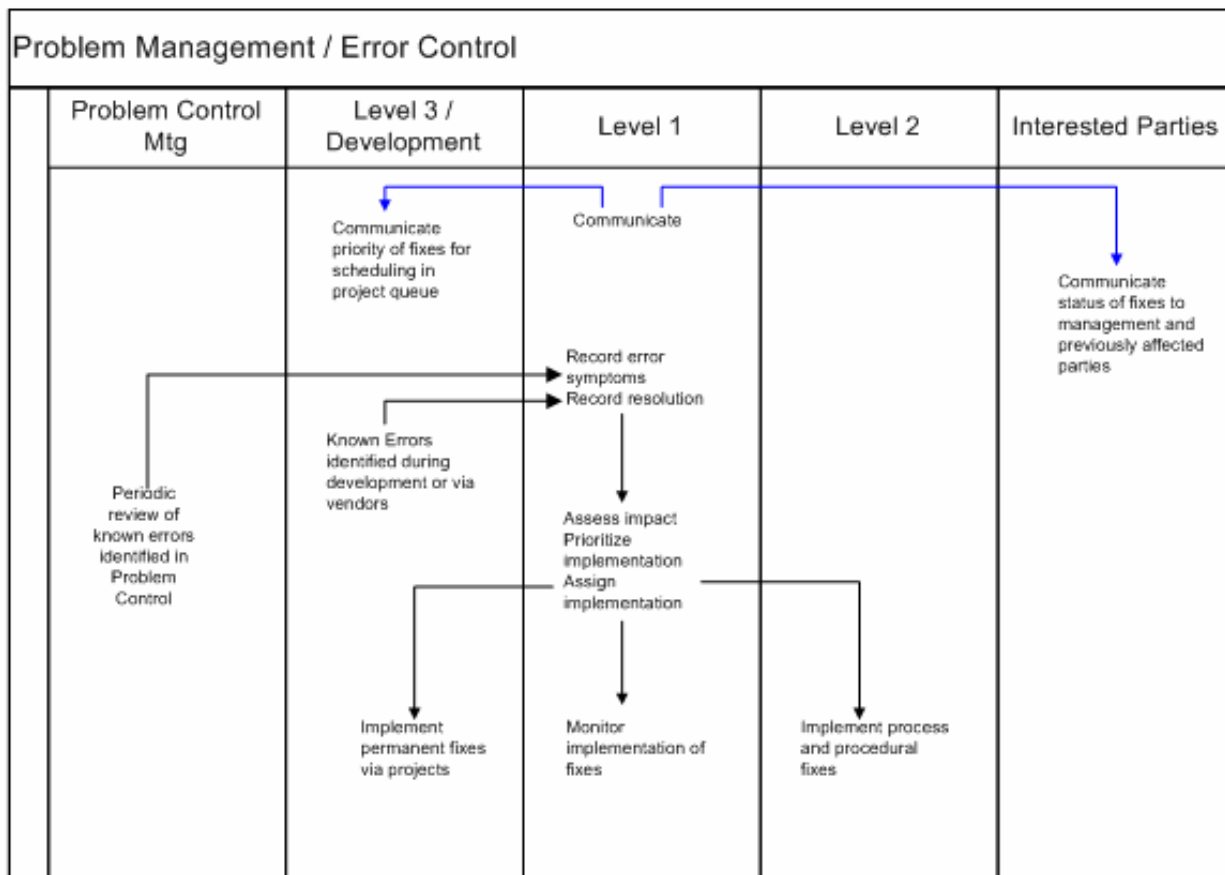
Error Control

Error control ensures that workarounds are documented and communicated to support personnel. It is also the liaison activity with development and engineering organizations that ensures known errors introduced through change management into the infrastructure are identified and communicated to support personnel. Further, error control seeks to influence the development/engineering organizations to implement permanent fixes to known errors. Table 3 shows the basic elements of the error control process. Figure 7 shows the basic error control process.

Table 3: Error Control Process Elements

Process Element	Description
Purpose	Communicate workarounds for known errors and ensure that these errors are fixed by development and engineering teams
Owner	Level 2 support team
Input	Problems whose root cause has been identified Known errors being implemented via change management process
Output	Documented workarounds to different support groups for known errors Prioritized list of projects to engineering/design groups to correct known errors
Measurements	Quantity of known errors Quantity of incidents caused by known errors Quantity of projects funded/implemented to fix known errors Sum of cost of all projects to fix known errors

Figure 7: Error Control Process Model



Communications

Communication tends to take two forms during IM and PM. Status messages of varying content are provided to different groups and/or individuals based on strict schedules and templates. Request/notification messages, which require the receiver to take action, usually have little content above and beyond the actual request, such as a reference to an incident ticket number, conference bridge telephone number, or other call-back number.

Technology

Most companies utilize text paging systems to facilitate automated communications during IM due to the time-critical nature of incident resolution. They can provide a significant level of detail via the text message, and have the ability to facilitate both kinds of communication: provide status and request assistance. Two-way text pagers (such as “Blackberry”) are used in many environments. Best-in-class IM processes integrate text pagers with the IM software to ensure that automatic status and request/notification messages are sent according to predetermined schedules.

Automated

Most companies rely on the automated communications capabilities of their IM software to automatically provide status and request/notification messages. The status messages from these systems are commonly generated from data entered into fields of the incident ticket itself. Therefore, these status messages are often cryptic and incomplete due to the fact that the fields used to construct the automated message may not be regularly updated with timely information or are automatically populated by monitoring tools with error message jargon.

Best-in-class organizations tend to rely on this automated communications capability to ensure that request/notification messages are sent in a regimented fashion to support escalation. They rely on the automated status messages for low-severity incidents, which are usually generated when a ticket is opened and closed. High-severity incidents are usually augmented with manually generated messages.

Manual

Manually generated communications are most often used to convey status during high-severity incidents. Manual communications tends to occur via e-mail and telephone because the message can be customized and different (non-email/pager) formats can be used. Templates or checklists are used to ensure that key information is conveyed.

Escalation

Escalation mechanisms assist in the timely resolution of an incident by increasing the staff capability, level of effort, and priority dedicated to resolving the incident. Best-in-class organizations have well-defined escalation paths with timeframes and responsibilities clearly defined at each step. They use the IM tool to automatically transfer responsibility to increasingly higher levels of support according to severity and timeframe.

Timeframes and responsibilities for escalation vary widely by organization, industry, and severity level of the problem. Best-in-class organizations negotiate with end users to determine appropriate timeframes and escalation responsibilities. The results of these negotiations are implemented in service level agreements, automated tools, check lists, templates, and individual performance reviews/objectives.

Functional Escalation

Functional escalation is the transfer of an incident to a higher-level support group when knowledge or expertise is lacking or when agreed on time intervals elapse. Best-in-class organizations define a matrix of severity levels based upon impact to the business, resolution timeframes and intervals at which the incident should be escalated to a higher-level group. Table 4 is an example of such a matrix.

Table 4: Escalation Matrix

Severity Level	Description	Resolution Goal	Entry Level	First Escalation	Second Escalation	Third Escalation
1	50+ Users unable to transact business	2 hrs	Level 1 Support	0 Minutes Level 2 Support	30 Minutes Level 3 Support	30 Minutes 1 st Manager War-room
2	10-49 Users unable to transact business	4 hrs	Level 1 Support	0 Minutes Level 2 Support	60 Minutes Level 3 Support	60 Minutes 1 st Manager War-room
3	1-9 Users unable to transact business	8hrs	Level 1 Support	30 Minutes Level 2 Support	120 Minutes Level 3 Support	120 Minutes 1 st Manager

In most organizations, level 1 and level 2 support groups are dedicated to providing support. Level 3 support, however, is usually provided by groups that have planning and design responsibilities. Thus, carefully planning how responsibility will be functionally escalated to level 3 is critical. Best-in-class organizations typically specify an “on-call” pager for each major technology group. The manager for each technology group is responsible for creating an on-call schedule and ensuring that the on-call pager is carried at all times. Further, each technology group must designate a managerial (hierarchical) escalation path. Typically, the line manager for the level 3 group is the first manager in the escalation path.

Hierarchical Escalation

To ensure that appropriate priority and resources are being allocated to resolve an incident before resolution timeframes are exceeded, hierarchical escalation involves management in the process. Hierarchical escalation can occur at any support level. In Table 4, hierarchical escalation occurs with the 3rd escalation for all problem severity levels.

Escalation to management occurs automatically in best-in-class organizations according to a predetermined schedule based on severity of the problem. When an escalation occurs, the targeted manager is expected to actively manage resolution of the problem, and becomes the single point of contact for status messages.

Cross-Group Escalation

A message to another on-call support group requesting assistance may be sent via text pager and/or telephone call, but responsibility for resolution of the problem is not transferred. The receiving support group is expected to render assistance to the caller. A pre-established audio-conference bridge is used to conduct the conference, which may grow to multiple support groups throughout the course of the incident.

War-room Escalation

The purpose of a war-room is to bring together key parties involved in restoring service to establish a plan of action and ensure that an appropriate cross-disciplinary response is being fielded. A war-room notification is automatically sent when it is likely that a high-severity problem will not be resolved within the target timeframe.

Also, war-rooms may be manually called by any manager to whom incident responsibility has been transferred. The war-room is usually a centrally located conference room with a speaker phone, whiteboard, network connections and a pre-established, toll-free audio-conference number. The war-room notification is sent to a predetermined list of attendees and the agenda for the war-room meeting is owned by manager to whom responsibility has been escalated. All people notified are expected to attend the meeting in person or via audio-conference.

The initial war-room meeting establishes the following:

- ▶ What the incident is
- ▶ Who is affected
- ▶ What the ETA for resolution is
- ▶ Who is working on the incident
- ▶ What is preventing the incident from being resolved
- ▶ Potential for not meeting resolution timeframes
- ▶ The plan of action
- ▶ Additional resources needed
- ▶ War-room meeting schedule until the incident is resolved
- ▶ Any exceptional communication that is required regarding the status of the incident

Reporting

Best-in-class organizations use reporting to control their IM processes, initiate continual improvement, and communicate performance against service level objectives to business unit customers.

Process Control

Reports generated from the IM system should be used to control the IM and PM processes. These reports allow managers to guide the effort of staff involved in incident and problem management. Examples of process control reports include:

- ▶ Quantity of incident tickets presently open by severity, longest elapsed time, responsibility group
- ▶ Quantity of problem tickets presently open (without root-cause identified)

Continuous Improvement

Reports such as the following can help identify weaknesses in the IM and PM process or even weaknesses in IT infrastructure itself.

- ▶ Mean time tickets were assigned to each level
- ▶ Quantity of tickets escalated and resolved by each support level

Service Level Performance

A vital set of reports allow IT to communicate to its customer's, performance against service level targets.

- ▶ Percentage of incidents resolved by resolution time-frame target
- ▶ Mean time to restore service

Conclusion

While most organizations develop process and procedures around incident management, many fail to do the same for problem management. Often this is due to a lack of clear understanding of the characteristics of the two activities. Incident management is the simplest activity to understand because it involves putting structure around the response to service interruptions. Because the “squeaky wheel always gets greased,” incident management discipline tends to develop quickly. However, there is often less insistence to develop discipline around problem management.

Problem management more closely resembles the management of a portfolio of projects, where each project’s objective is to determine the root cause of a problem. Incidents are often the first indicator of a problem and once revealed by an incident, the organization must have process and procedures to get to the root cause. To continue with the portfolio analogy, an organization that has addressed problem management will develop criteria to determine which problems should be investigated to determine root cause, in much the same way that an organization will have some decision criteria for committing to a new project. Problems that are not investigated will continue to be tracked and monitored for future investigation. When root cause is determined and a solution developed, the organization will track the progress of implementing the solution.

INS encourages organizations to consider the ITIL-based framework presented in this whitepaper. The development of process, procedures, tools and organizational changes to implement an incident and problem management system can be a tremendous undertaking. In many cases, changes to organizational culture are required to fully adopt the problem management elements. Organizations may find that this undertaking is more easily accomplished through the use of a consultancy, such as INS, to augment the organizations own efforts.

About INS

INS (International Network Services Inc.) is a leading global provider of vendor-independent network consulting and security services. We offer a full range of consulting services to help companies build, optimize, manage, and secure their network infrastructures to enable their business initiatives and achieve a sustainable operating advantage. INS is a recognized leader in business-critical, multivendor network consulting, having helped more than 75% of the Fortune 500 and delivered more than 15,000 engagements over the past decade. Headquartered in Santa Clara, CA, INS has regional offices throughout the United States and Europe. For additional information, please contact INS at 1-888-767-2788 in the U.S., 44 (0) 1628 503000 in Europe, or 1-650-318-1020 worldwide, or visit www.ins.com.

Copyright © 2003, International Network Services Inc.

This is an unpublished work protected under the copyright laws.
All trademarks and registered trademarks are properties of their respective holders.
All rights reserved.