Problem Management

Contents

Introduction

Networks provide the highest levels of availability and performance when operations are predictable and repeatable. A key challenge to network support organizations is the establishment of a process that can detect, diagnose, and resolve existing or potential problems that prevent the achievement of the desired availability and performance goals. Application of the correct mix of people, processes, and tools guided by accepted industry best practices, as defined in sources such as the IT Infrastructure Library (ITIL), greatly contributes to predictable and repeatable levels of operational performance. However, implementing this best practice and the supporting management infrastructure can often present challenges such as the following:

- Making sure that resources have the right skills and can commit sufficient time to problem management
- Making sure that processes are in place to identify, classify, diagnose, and control current and potential problems
- Making sure that the right set of management data is collected and that tools are available to support the problem management process
- Making sure that the information discovered and the changes recommended by problem management are distributed to the rest of the operational environment

The Cisco® commitment to driving operational excellence into the management of Information and Communication Technology (ICT) has been realized through a number of channels including the development of a wide portfolio of management tools, support to other management tool vendors through development partnerships, and efforts in the standards arena.

Support to customers has been realized through the distribution of best practice guidance and professional services that assist customers with implementation of both tools and key best practices into their operational environment.

As part of this ongoing commitment to operational excellence, this paper introduces the key concepts that underpin the development of an effective problem management process, the issues likely to be encountered when implementing such a process, and the options that can help mitigate these issues.

Overview

Goal of Problem Management

The goal of problem management is to minimize both the frequency and impact to the business of conditions that reduce the availability or performance of the network infrastructure. Problem management in a networking context is the structured analysis of both network performance and network outages to establish the root causes of issues that, when addressed, improve network availability and performance.

Effective problem management requires a rich set of management information to support diagnosis of root-cause conditions currently affecting the network and as the basis for trend analysis that will diagnose trends in network usage, which, if not addressed, could threaten availability or performance goals in the future. The correct selection of tools to provide this information, the correct training of staff in the use of these tools, and the orchestration of diagnosis, resolution, and knowledge management through a robust business process is critical if problem management is to realize benefits.

Components of Problem Management

Problem management is a business function composed of people, processes, and tools organized and chartered to resolve customer problems (Figure 1). A problem is defined as a cause, or potential cause, of one or more incidents that have affected the availability or performance of the network infrastructure.

Figure 1. Problem Management Core

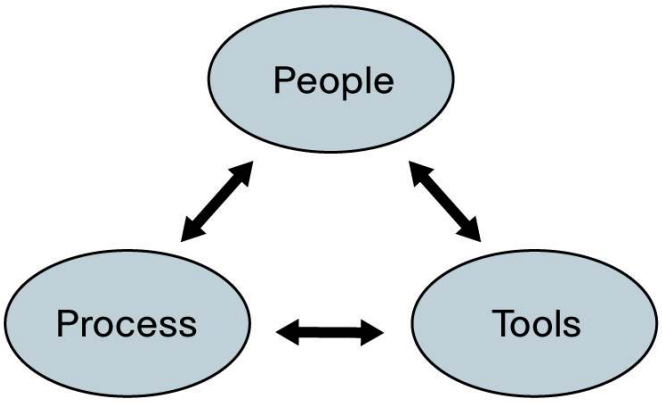

Problems are managed through the problem management process from their detection through their elimination from the network. Throughout the process, information is collected that is useful in addressing new incidents, as well as for continued governance of the process itself. The primary objectives of problem management are to:

- Detect problems or problem indicators before users are affected
- Identify root causes of problems and provide workarounds for use by tier I support
- Prevent incidents by addressing their underlying causes
- Minimize the impact of incidents that cannot be prevented

Challenges to Effective Problem Management

All network operators invest a significant amount of manpower and financial resources into maintaining and servicing the network to meet the expectations of their users. With the integration of advanced technologies and applications, services have become more complex, and the rate of change in the infrastructure has increased along with users' expectations of performance and availability.

These challenges have placed increasing strains on the ability of IT organizations to scale support to encompass all the activities required to support this dynamic mix of infrastructures.

It is all too common for an organization to become fully engaged to the point that most resources are committed to reacting to incidents and quickly restoring service, leaving little time for structured analysis of these incidents to determine the underlying root-cause conditions that are at their source.

The management plane of the network infrastructure generates a very rich set of data to be mined when diagnosing problems. One of the key challenges is to make sure that the volume of management data itself does not become an issue that prevents effective problem management. The lack of alignment of management information and the configuration of the management tools that collect and present this information to address the needs of effective problem diagnosis and resolution is another implementation challenge to be addressed.

The problem diagnosis process itself can involve multiple stakeholders whose activities need to be coordinated and tracked to make sure that the cost of investigation itself does not become an issue and that effective resolutions to problems are actually discovered. Selecting and implementing the correct tool to manage these activities is a key challenge when implementing a problem management process. Facilitating the effective reuse of the knowledge discovered during

the problem management process by other parts of the support organization is another challenge to be addressed.

The final challenge in realizing the benefits of a problem management process comes when attempting to implement the remedial changes identified through the process. Effective remedies may involve modifying organizational behavior or making capital investments in the network infrastructure. Securing support for remedial changes and ensuring successful implementation is the critical element in the realization of an effective problem management process.

Difference between Problem and Incident Management

An incident is an instance of a condition that has affected the availability or performance of the network infrastructure. The focus in managing incidents is rapid diagnosis and rapid restoration of affected network services. While incident management may recognize that the incident is the result of an underlying problem, the need for rapid resolution will usually demand a solution that works around the underlying problem rather than resolving it.

Problem management is the systematic review of past incidents and the trends in network usage to identify the changes required to permanently remove these underlying problems from the infrastructure and prevent future incidents. The focus in problem management is on in-depth investigation and fundamental change to the network infrastructure. As such, timelines for diagnosis and resolution are considerably longer than for incident management. For example, an incident might be resolved by making a minor configuration change to the network. However, the underlying problem may require an adjustment to the network design.

Although operating over different timescales, problem and incident management are closely related, and benefits are derived when commonly agreed categorizations for impact and priority are applied and toolsets permit both functions to share information.

Benefits of Problem Management

The effective execution of problem management will help the support organization proactively establish the underlying causes of problems or potential problems. Benefits of implementing problem management include:

- Improved service quality
- Incident volume reduction
- Permanent solutions
- Reduced cost of reactive support
- More timely resolution of outages
- Higher productivity of both the business and IT

Process

Problem Management Process

The problem management process can be viewed as supporting four broad functional areas (Figure 2). These are:

- Problem Control

Capture the details of a problem, validate its existence in the network, determine its priority relative to others, and assign it to a group for investigation.

- Error Control

Capture details of a known error and distribute information to other support groups, define steps to control resolution of the error, and monitor the implementation of these resolution steps.

- Proactive Problem Management

Perform the activities necessary to minimize opportunities for incidents to occur in the infrastructure by identifying and eliminating underlying problems that cause them.

- Major Problem Review

Perform a “lessons learned” exercise as part of a commitment to continuous improvement of the problem management process itself.

Figure 2. Problem Management Process High-Level View (Source: ITIL Service Operation – OGC)

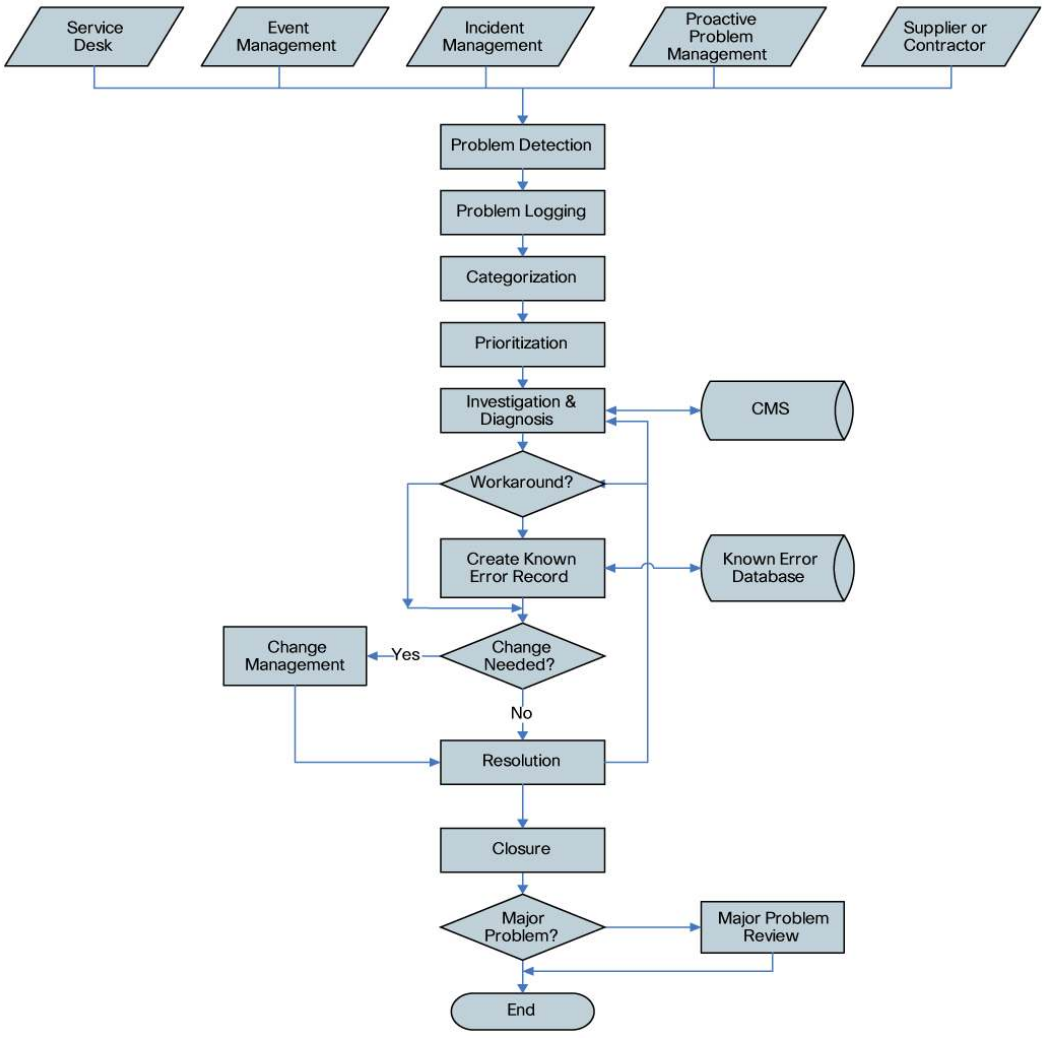

The activities performed in each of these areas and the characteristics of successful operation are discussed in the following sections.

Problem Control

Problem control is the systematic analysis of the incidents that have affected the availability and performance of the network in order to identify the problems that need to be investigated, the

priorities to be assigned in the problem investigation process, and diagnosis of the root cause. Problem control is itself composed of the following steps:

- Problem Detection

The identification of the existence of a problem is the first step in the problem management process. Problems can be identified by a number of different sources:

  - Service desk identifies a potential problem through correlating information or attributes of multiple incidents or a single type of incident experienced multiple times.
  - Technical teams, such as software developers, analyze an incident and determine that an underlying problem exists or is likely to exist.
  - Vendors and other third-party updates/notifications reveal problems.
  - Automated detection of an infrastructure fault is made through network management tools.
  - Proactive problem management activities identify a condition or vulnerability that requires further investigation.

A support organization must have sufficient time to perform the types of analysis required to identify problem conditions and to validate the presence of problems being reported. It is also important that the expectations of the problem management team, in terms of supporting information required from other teams when reporting problems and the supported reporting methods, are clearly communicated throughout the organization.

- Problem Logging

Once the existence of a problem has been validated, it is important that details of the problem itself are documented (Figure 3). A problem record, as it is termed, captures this initial set of information and is then used to document additional information that is generated in the later phases of the problem management process.

The choice of tool to store and manage problem records is a critical success factor in implementing an effective problem management process. This is further discussed in the tools section of this paper. However, a tool is only as effective as the people who use it. Organizations who take the time to make sure that the quality of the information being captured is sufficient to support problem management—as well as to other key process areas needed to drive higher levels of availability into the network—tend to be the most successful.

Figure 3. Problem Logging Process

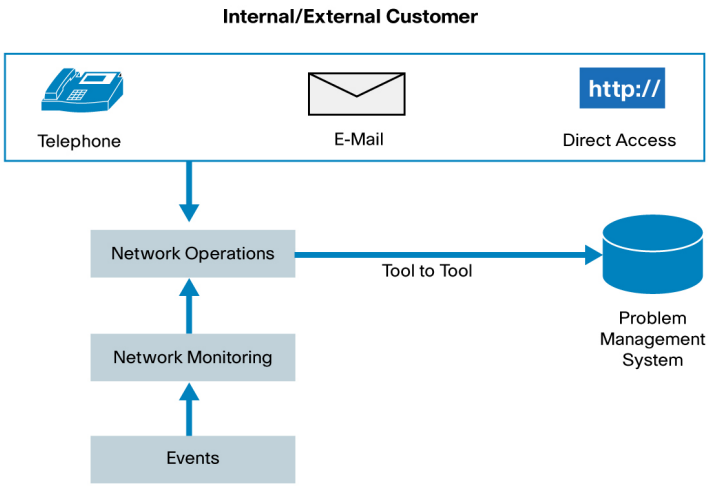

- Problem Classification

Problem classification (Figure 4) is the process of assessing the impact a problem is having on the business and determining who should be investigating it. Problem records are updated with an attribute that defines the category of the event, for example, hardware, software, and so on, and is used to allocate the correct technical resource to the investigation of the problem.

Figure 4. Sample Problem Classifications

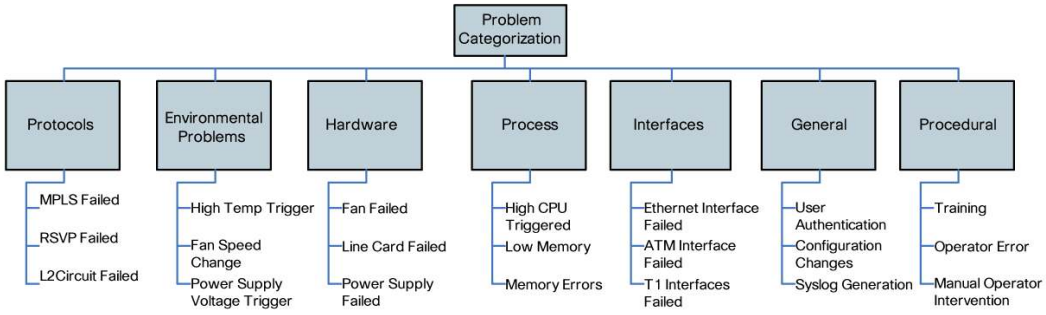

An attribute is added to the problem record that defines the impact of the problem and allows the scope of the investigation to be determined.

- Problem Prioritization

Problem prioritization (Figure 5) is the process of determining when a problem should be investigated. An attribute is added to the problem record that defines the urgency of the problem. Urgency is an assessment as to whether the problem must be investigated immediately or can be scheduled for investigation at a later date.

After impact and urgency are determined, an attribute is also added that designates the priority of the problem. As a team may be managing multiple problem investigations concurrently, priority and urgency are used to determine the priority order in which investigations occur (Table 1).

Figure 5. Decision Tree for Determining Priority (Source: IT Problem Management by Gary Walker)

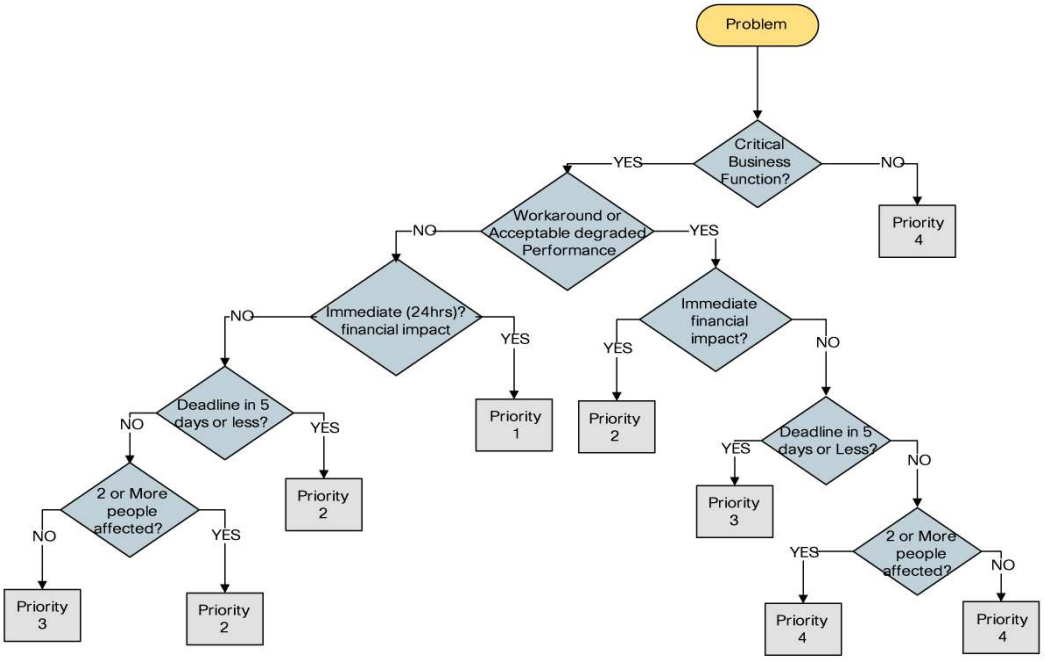

Table 1. Problem Management Priority Matrix

Priority Level Number	Priority Definition
1–Critical	Failure of a component where one or more people cannot perform critical business functions. Failure to complete this business function within 24 hours will have a negative financial impact on the company. No workaround is available, and a degraded mode of operation is not available or not acceptable.
2–Urgent	Any of the following conditions is true: Failure of a component where one or more people cannot perform a critical business function. Failure to complete this business function within 24 hours will have a negative financial impact on the company. A workaround is available or a degraded mode of operation is available and acceptable. or Failure of a common component where two or more people cannot perform a critical business function. This failure will not have an immediate financial impact and there is no deadline within 5 days. No workaround is available and a degraded mode of operation is not available. or Failure of a common component where one or more people cannot perform a critical business function and are at risk of not meeting a deadline for the critical business function in 5 days or less. There is no immediate financial impact. No workaround is available.
3–Important	Any of the following conditions is true: Failure of a component where one person cannot perform a critical business function. There is no workaround available. Failure of the business function will not have an immediate financial impact and there is no deadline of 5 days or less at risk. or Failure of a component where one or more people cannot perform a critical business function. None of the affected people has an immediate negative financial impact. One or more of the affected people has a deadline of 5 days or less. A workaround is available or a degraded mode of operation is available and acceptable. or Failure of a component where two or more people cannot perform a critical business function. None of the affected people has an immediate financial impact and none has a deadline of 5 days or less. A workaround is available or a degraded mode of operation is available and acceptable.
4–Low	Either of the following conditions is true: Failure of a component where one person cannot perform a critical business function. This failure will not have an immediate financial impact and there is no deadline within 5 days or less at risk. A workaround is available or a degraded mode of operation is available and acceptable. or Failure of a component that affects a noncore business function

Monitor	No business impact. Does not affect a core business function, for example, information requests and scheduled events.

Successful organizations establish well-defined schemes for these attributes and reuse them across all the processes that are associated with managing the availability of the network. Communication of the assigned priority and urgency to affected users is also a key to effective problem management, as is the process of enforcing these values in the face of pressure from other parts of the business.

The ability to alter priorities based on changes in business conditions or a technical driver should exist but should be exercised with due diligence. If all problems are assigned high priority and high urgency then these attributes become meaningless and the problem management process itself becomes ineffective.

Successful organizations also are strongly focused on using tools that permit linkages between incidents, problem records, and configuration management data to be established. Such linkages allow the impact of a problem to be fully qualified and priorities to be defined that are an accurate reflection of the impact on the business.

- Problem Investigation and Diagnosis

The goal of the problem investigation phase (Figure 6) is to identify the root cause of the problem so that appropriate remedial steps can be put into action. Problem investigators can utilize a variety of network management and other tools and may be required to work closely with other support groups and users in order to establish the root cause. Other information, such as procedural or product or design documentation, may also be used in the investigative process.

Whether the root cause of a problem is related to equipment failure, configuration error, network design, patterns of usage, or process failure, the investigation phase should result in an update to the problem record that reclassifies the problem as a known error. This reclassification is also accompanied by a recommended workaround that will be used to manage incidences of the problem until the changes in the network infrastructure or wider organization that will permanently remove it have been achieved.

Figure 6. Problem Investigation Process (Source: IT Problem Management by Gary Walker)

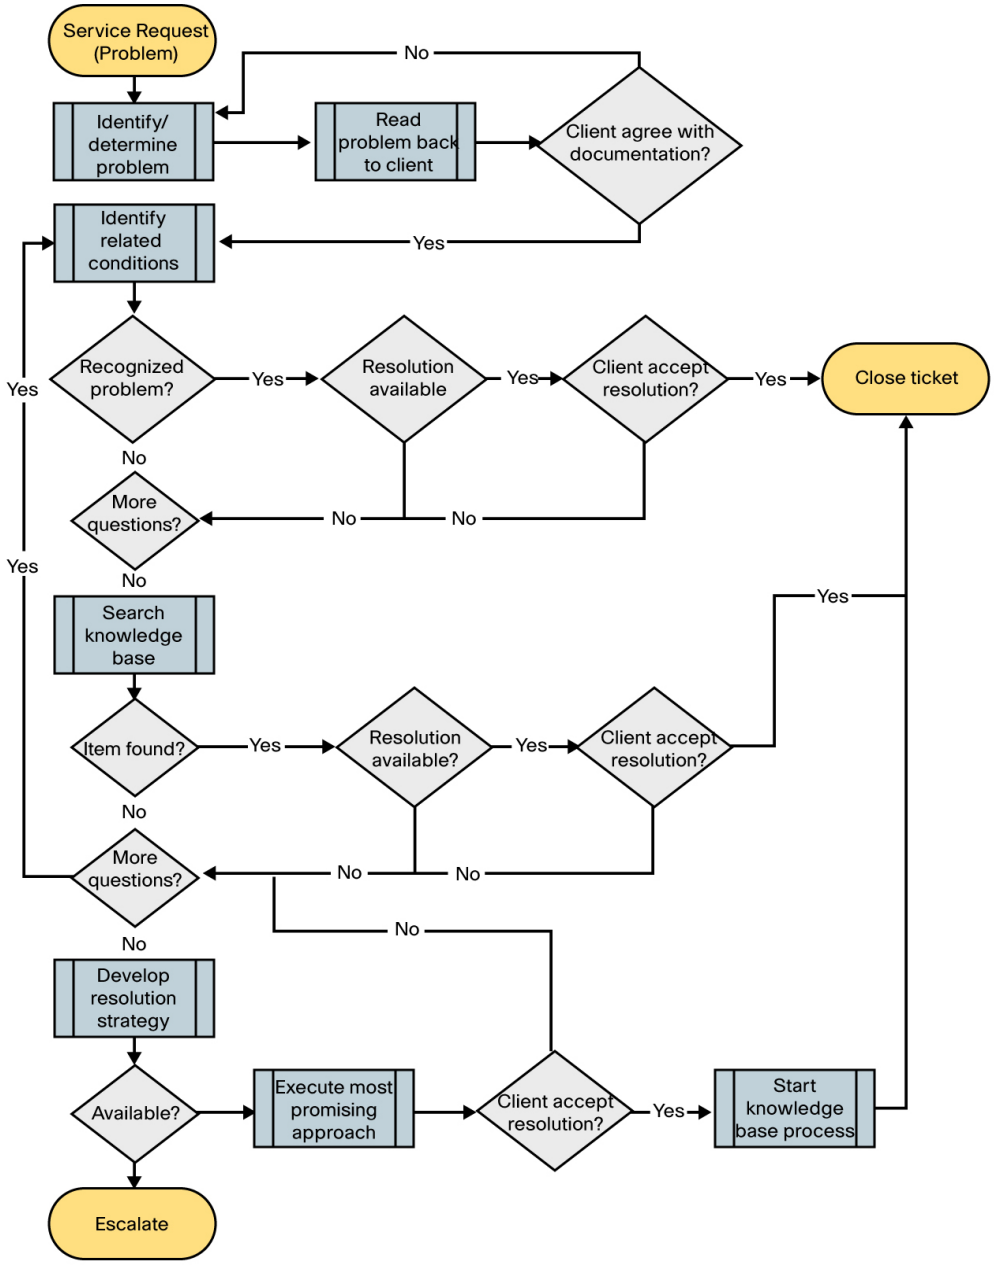

Successful organizations are those who commit sufficient resources to investigate problems and work successfully with other parts of the organization to establish the root cause. A successful organization is also characterized by the extent to which known errors and workarounds are distributed to other support groups and used to manage future consequential incidents until the problem can be corrected. A successful organization has established mechanisms for collecting the information required for problem diagnosis and procedures for handling any disruptive testing that may be required as these often affect the length of the investigation phase. These organizations typically also have good visibility of information about change in the network, as this is often the root cause of identified problems.

Error Control

The error control phase of the problem management process is concerned with the management of known errors until they are successfully resolved. Error control is composed of the following activities:

- Error Identification and Recording

The error control process starts as root causes of problems are identified and workarounds are created and documented. Problems are thus transformed into known errors and documented in a known error record. Known error records can then be used by support personnel to aid in the handling of incidents.

Successful organizations have problem control processes that can supply known errors from both live and development environments. They also have processes in place to generate known error records for supplier-reported problems.

- Error Assessment

The planning required for successfully implementing remedial actions to resolve a known error occurs in the error assessment phase. This phase of the process will take inputs from all the relevant contributors to the resolution, be they internal resources or external suppliers, and define the changes required. Where changes require modifications to the network infrastructure, integration with the change management process occurs through the raising of change requests.

Successful organizations are those that have strong contractual relationships with external suppliers so that resolution of known errors is delivered in timescales that meet the needs of the business. Successful organizations also have internal controls through which supplier performance in the resolution of known errors can be tracked and enhanced if failure to meet expectations is encountered.

Successful organizations are also characterized by their emphasis on the regression testing of the resolution to known errors to make sure that changes are implemented with minimal disruption.

- Error Resolution Recording

In addition to the information captured in the problem control phase, the known error record is used to capture information that defines the steps to be implemented during resolution and information that permits progress towards resolution to be tracked.

Successful organizations store known error reports in a known error database to make the information available to other support groups while resolution is being implemented or to document those errors for which resolution is deemed to be not cost effective

As with other aspects of information recording, successful organizations implement standardization in the presentation of information in the known error reports and have quality assurance measures in place to make sure that the information in the database is being well maintained. Please refer to the tools section for further information on features of systems to manage known error reports.

- Problem/Error Resolution Monitoring

The continuous monitoring of the activities that are required for detecting problems and resolving known errors is necessary to make sure that the problem management process as a whole is being effective.

The definition of measures of performance that can be used to track the effectiveness of the process against service-level agreements (SLAs) and the regular reporting of these to management is essential to help ensure viable operation. Typical metrics measure trends in the volume of problems and known errors, rates of detection or resolution, and costs associated with the operation of the process.

Successful organizations define service-level agreements and reporting mechanisms around the resolution process to make sure that problems are detected and resolutions are being implemented in a timely and cost-effective manner. Successful organizations also have internal controls that allow for the priorities of problem management activities to be reviewed and, if necessary, adjusted to reflect changing conditions.

- Error Resolution Closure

Closure of a known error record occurs once the steps outlined in the plan for resolution have been successfully implemented. Any correlated problem or incident records that are related to the known error are also closed in this phase of the process.

Organizations exhibiting best practice will typically assign a status of “closed pending review” of the execution of the remedial steps. This interim status allows the final confirmation that the known error has been removed from the environment. Successful organizations also have internal controls that help ensure that known errors are moved to a final “closed” status and do not remain in the “closed pending” status.

Proactive Problem Management

Proactive problem management is the process of reviewing the operation of the network infrastructure from a number of different perspectives to detect the presence of problems before they give rise to incidences of network unavailability or degraded operation.

Organizations with a commitment to proactive problem management utilize data from the tools that manage the network, design material, test reports, user experiences, process reviews, and external suppliers to determine conditions that could affect availability and implement the steps necessary to resolve them.

Typical activities might include the monitoring of network usage patterns or component failure rates to identify trends, which, if left unchecked, may lead to incidences of network unavailability or degraded operation.

Successful organizations recognize the importance of this activity and are prepared to invest the necessary resources to make the process effective. Successful organizations are those that follow through on implementation of the recommendations from proactive problem management activities and do not allow them to be marginalized by other priorities.

Major Problem Review

Major problem reviews are an invaluable tool for organizations committed to a culture of continuous process improvement, as they allow the lessons learned in execution of the process to be incorporated into subsequent operations. The review process should focus on what went well, what went badly, what could have been done differently, and what needs to be in place to avoid the issues in the future. Successful organizations are those that conduct reviews that incorporate the views of multiple stakeholders involved in managing the problem to resolution, act on the lessons learned, and can measure the impact through the metrics that are used to report on the

effectiveness of the problem management process.

People

Staffing and Support Structure

Organizations differ in their approach to staffing and support structures for infrastructure networking environments based on technology and network life-cycle requirements. Smaller organizations are often able to implement a cradle-to-grave approach where one group handles planning, design, implementation, and operation for an individual technology or a small set of technologies. Larger organizations will generally require a center of excellence support structure or a converged approach. This center of excellence approach separates technologies and lifecycle functions into separate managed entities. A converged approach may separate lifecycle functions, but seeks to combine different technologies that are used to support an individual solution. See Table 2.

Table 2. Benefits and Issues with Different Organizational Structure Types

Organizational Approach	Benefits	Potential issues
Technology-based cradle-to-grave	• Generally good teamwork and communication • Higher levels of expertise through on-the-job training • Easier to manage	• Multiple technology groups may compete or overlap in converging network environments • Difficult to scale
Center of excellence	• Separation of duties • Can be more cost-effective with skills only required within the silo • Less opportunity for competition • Scales well	• Communication and teamwork sometimes difficult • Knowledge silos may develop with less technology sharing • Requires strong process management
Converged	• Different technology teams work together to provide the best solution and support • No finger pointing between technology groups • Cost effective	• Commitment to aggressive cross-training • Merging groups with different cultures, personalities may be difficult

In infrastructure network environments, a hybrid or converged approach may work best. This helps avoid any competition in solutions between server and network groups and will help avoid finger pointing between different support groups. A converged approach may be a large step for some environments and is recommended only after careful consideration and planning within the individual organization. Any change to organizational structures should not be attempted unless it has management support and is properly communicated and promoted throughout the organization.

The staffing levels and organizational structure required for each of the above approaches will vary widely, but considerations should be given to:

- Overall number of incidents
- Number of unique network technologies
- Number of equipment vendors
- Number of service suppliers/partners
- Maturity of the network infrastructure
- Maturity of operational processes
- Complexity of network architecture

- Number of systems and applications being used
- Operational level agreements or service level agreements
- Expertise of existing staff
- Organizational culture
- Geographical responsibilities

The recommended method of determining staffing levels for problem management is to build a staffing plan that identifies roles and responsibilities related to network lifecycle functions and includes a skills matrix for current and planned staff based on technology skills and lifecycle process capabilities.

Roles and Responsibilities

Table 3 identifies the specific roles and responsibilities of the problem manager.

Table 3. Problem Manager Responsibilities

Problem Manager Responsibilities	Recommended Background/Qualifications
• May be coupled with the trouble triage of service affecting incidents and mission-critical incidents • Maintain log of major isolation, restoration, and resolution steps • Convert problems into known errors • Raise requests for changes • Identify, document, and record workarounds • Work with various support staff to help ensure resolution and closure of tickets • Perform analysis of major incidents to identify chronic troubles • Conduct Pareto analysis of incidents to categorize troubles as product, processes, or skills • Interpret near real-time and historical network or IT performance indicators to recognize potential problems and act on them before they cause incidents • Identify countermeasures to drive quantifiable improvement • Assemble virtual teams to address countermeasures • Engage vendors and suppliers in countermeasure team • Facilitate and moderate conference calls with countermeasure team to help ensure thorough root-cause analysis and creation of viable countermeasures • Analyze countermeasures and drive implementation of high benefit items • Manage log of countermeasures • Share lessons learned with service desk and operations • Provide proactive analyses of trend data to identify potential outages and failures	• Ability to direct a cross-organizational team • Strong problem identification, isolation, and resolution skills for complex network problems • Seven years of experience in working in a medium to large enterprise or service provider environment • Cisco Certified Network Professional (CCNP®) or preferably Cisco Certified Internetworking Expert (CCIE®) qualification or equivalent • Working knowledge of diagnostic tools and protocols • Working knowledge of change management procedures and configuration templates • Working knowledge of network documentation – software, physical and logical elements, archival and retrieval • In-depth understanding of network hierarchy, architecture, protocols, and overall network design • Strong analytical, interpersonal, and communications skills • Working knowledge and understanding of vendor's tools • Working knowledge of performance and capacity management tools • Knowledge of security management tools and procedures • Familiar with project management tools and techniques

Levels of Expertise

The levels of expertise within an organization can often affect the staffing levels required, the quality of the solution, and the time required to resolve problems. A skills matrix identifies training and expertise gaps in the organization, and helps to identify additional training needs. Skill values can sometimes be based on industry certification capabilities, if available. Skills matrices generally utilize a 1–5 scoring mechanism, given in Table 4.

Table 4. Skills Matrix Rating

Skill level	Description
1	Beginning, no familiarity with problem management

2	Can manage basic problem resolution given job-aids and coaching
3	Can analyze most problems and corresponding root-cause analysis (RCA) and proposed countermeasures
4	Can analyze all problems and corresponding RCA and proposed countermeasures
5	Can plan and manage the implementation of countermeasures to prevent problem recurrence and improve overall service performance. Independently directs cross organizational virtual teams.

This skills matrix rating is then used to evaluate where the problem management organization is in relation to the product and technology mix. An organization may also use this skill matrix in relation to its current organizational structure to identify gaps in expertise as depicted in Table 5.

Table 5. Skills Gap Analysis Matrix

Functional Role	Problem Manager
Managing a virtual team	3
RCA assessment	2
Analysis of countermeasures	2
Reach consensus with stakeholders on countermeasures	1
Project manage implementation of countermeasures	No skill

Training

To build problem management expertise, the organization should adopt training strategies that include formalized training courses, on-the-job experience, and periodic skills evaluation. This method works best when experience or lab capabilities exist (immediately following training) and strong mentoring is immediately available. Cisco provides evaluation/certification opportunities within the Cisco Certified Network Professional program on Cisco.com.

Train the team in the processes. Make sure the team members understand every step of the process. Make sure they understand the inputs and outputs of each step. Also, make sure the teams understand the logic behind each step in the process. People tend to be more effective when they understand why they are doing something, not just what has to be done. Train them how to work together. Provide leadership through communicating the vision of how it will work. Set explicit expectations on the level of cooperation you expect and the level that is required to make the processes work.

Communication and Teamwork

It is essential that a high level of communication and teamwork is in place and actively encouraged to achieve highly available network/service environments. Often, especially in large organizations, the delivery of a service must traverse network infrastructures affecting multiple support groups, departments, and business units. Without a “one team” perspective and commitment to open communications, troubleshooting and problem resolution can be dramatically affected, jeopardizing long-term network health and the ability to meet service-level agreements with customers.

Successful organizations supporting highly available service environments have clear levels of understanding and communication between the various support groups, devoid of “blame cultures”. They have a proactive behavior policy model, versus reactive. In some cases, to encourage such high levels of cooperation, organizations have embraced the concept of “virtual teams” with regular meetings to facilitate mutual knowledge sharing and flow of information. Rotating jobs between support team members is also seen as a best practice approach to encouraging teamwork and better communication as well as improving skills and broadening

perspectives.

Measurement

As mentioned above, metrics are the key to any successful program and process. The establishment of measurable goals and objectives for the organization is a key step. Subsequent breaking of high-level organizational goals and objectives into smaller team-level goals and objectives so they can be cascaded to individual teams integrates these goals into the daily operation of the process.

Tools

Selecting a Problem Management System

The choice of an information system to support the business processes should be driven by an assessment of the ability of the system to bring consistency and economies of scale to the operation of the various components of the problem management process.

Key features that should be satisfied by an information system that supports problem management are presented below.

Functions to Be Supported

An effective problem management system must be capable of supporting a range of features that are required for successful execution of the problem management processes. Some of the main functions that a system must be capable of supporting are:

- Application administration
- Management of the problem control component of the problem management process
- Management of the error control component of the problem management process
- Management of proactive management component of the problem management process

Table 6 maps these broad functional criteria into specific features that any tool being considered should be capable of supporting.

Table 6. Problem Management System Functional Matrix

Process Area Functional Grouping	Subprocess	Tool Feature
Application administration	User administration	Ability to create users, modify user profiles, and remove users
		Role-based restriction of user access to product features
		Association of users to specific customer or group of customers
		Capture user contact details in a user profile
		Audit user activity and changes to problem and error records state
	Information management	Ability to create coding schemes for fields within the problem record
		Ability to define custom management reporting
	Integration	Ability to supply core functionality through a well-defined API
		Ability to conform with the requirements of the application integrations defined by internal IT
	Availability	Ability to operationalize the problem management solution in a redundant mode of operation
Problem control	Problem logging	Automated logging of problem records from an incident management system

		Manual creation of a problem record
		Ability to add customer location and contact details to problem records
		Attach details of affected network elements to the problem record
		Ability to associate correlated incident reports to a problem record
		Automated recording of date and time of problem record creation
		Ability for incident management system to gather current state of all problem records
	Problem classification/prioritization	Ability to assign initial values for urgency and priority of the problem investigation
		Ability to modify the values for urgency and priority if required
		Ability to associate the problem with a particular class of component
		Ability to associate the problem record to an instance of a network element
	Problem investigation and diagnosis	Ability to automatically route a problem record to the correct team of investigators
		Ability to assign a problem record to a particular investigator or group of investigators
		Ability for investigators to log investigative actions undertaken
		Ability to associate problem records with process or design issues as well as network elements
		Ability for investigators to record time spent in investigating problems
		Ability to record the cost of an investigation
		Ability to search problem records
		Ability to correlate problem reports to a common root-cause problem, making and breaking associations as required
		Ability to add attachments to the problem record
		Ability to allow recommended workaround to be added to the problem record
		Ability to use a variety of methods to notify investigators of changes to problem record
		Ability to transform a problem record into a known error
		Ability to close a problem record if no root cause can be determined
Error control	Error identification and recording	Ability to register problem records as known errors
		Ability to make known errors available to incident management systems
		Ability to define an SLA for resolution activities
	Error assessment	Ability to record the actions required to establish a suitable solution
		Ability to record performance of this subprocess
	Error resolution and recording	Ability to record the plan of action required to implement changes to the infrastructure
		Ability to establish a bidirectional interface to a change management system that can raise and manage change requests

	Problem/error resolution monitoring	Ability to provide a range of management reports, including but not limited to: • Total problems recorded in the reporting period • Percentage of problems resolved within SLA targets • Number and percentage of problems that breach SLA • Backlog of outstanding problems and the trend (static, reducing, increasing) • Average cost of handling a problem • Number of major problems (opened and closed and backlog) • Number of known errors added to the known error database (KEDB)
		Ability to notify managers when an SLA breach is likely to occur in resolving problems
	Error resolution closure	Ability to mark the status of a known error as closed
		Ability to remove closed errors from the KEDB
		Ability to notify problem and incident management systems of state changes
Proactive problem management	Problem detection	Ability to log the outcome of this process as problems to be tracked and resolved by the problem management process

Interfaces to Be Supported

To be effective in the support of a problem management process, the problem management system must be capable of interfacing to both users and applications. Outlined below are some considerations to be applied when reviewing the interface capabilities of potential candidates for a problem management system.

User Interface

The preferred interface should be a Web-based interface with comprehensive browser support. The interface should be well engineered, and transition between various steps in the use of the application should be intuitively presented.

The ability to customize the end-user experience should also be viewed as a requirement to be satisfied in reviewing user interfaces. The ability of the application to present inexperienced users with a lightweight, wizard-driven interface and more experienced users with detailed screens should be viewed as a plus. Also, the ability to restrict access to the resources in the system to those required to support an individual customer should be an important criterion.

The number of active elements on each page presented by the system also has a large impact on the quality of the user experience. A Web-based application that runs well in a high-speed LAN environment may run at a crawl when accessed through a WAN link. Evaluation of Web interfaces is necessary to test their speed across a range of access conditions to make sure that the right balance between usability and functionality is maintained.

Web interfaces typically make extensive use of Java technologies in order to support the end user. When reviewing an interface, it is important to make sure that supporting technologies such as the Java Runtime Environment versions align with the requirements of other Web-based applications that the users of the problem management system need to access.

Finally, support for the secure transport of user interactions with the problem management system should be considered a core requirement.

Interface to Incident Management

The incident management and problem management processes are closely related, and it is natural therefore to consider implementation of an interface between the toolsets that support both of these process areas.

The interface itself should be capable of permitting information about known errors and information in the knowledge base to be shared from the problem management system to the incident management system. In the other direction, the incident management system should have the ability to utilize the interface to create and manage problem tickets within the problem management system.

The nature of the interface will be governed by the nature of the two systems themselves. Toolsets purchased externally from the same supplier are likely to be tightly coupled from the outset, and all the functionality required to generate the information flows across the interface will be in place.

By contrast, problem and incident management systems that have been sourced from different suppliers—be they external vendors or different internal teams—will need to be integrated.

Integration challenges to be addressed include:

- The format of messages to be exchanged between both systems
- Synchronization of data presentation in areas such as classifications, severity, and other identifiers
- Choice of message transport
- Securing the message exchange
- Routing between primary and backup instances of incident and problem management systems
- Alignment with corporate application integration strategies
- Monitoring and reporting on the availability of interfaces

The absence of clearly defined standards for information exchange between incident and problem management systems is an issue and may limit the scope of any integration.

Interface to Configuration Management

The ability to share information related to configuration items and their relationships between configuration and problem management process areas greatly increases the effectiveness of problem investigations.

Such an interface should be capable of permitting the problem management system to interrogate the configuration management system and present information about configuration items in the returned data set.

Again, the absence of widely adopted standards for presentation of both search criteria and results presents the same challenges when integrating systems from different vendors. The adoption of simple data-level integration, while initially attractive, tends to become problematic when the configuration management database is a confederation of different data stores each with its own unique data model.

These issues are less apparent when the configuration and problem management systems are procured from a single source.

Interface to Change Management

An interface between the problem and change management systems permits change requests to be created and scheduled to remediate known errors.

Such an interface should be capable of permitting the problem management system to both create a request for change, as well as to modify the content of the change, or, if necessary, remove the request for change. As changes must be approved, the change management system must also be capable of sending asynchronous notifications of changes to the state of requests for change within the change management system. Since approximately 40 percent of network incidents result from change activity[1], the interface should be capable of populating the incident and problem tickets from the change record

As with the other interfaces, the nature of the integration challenges to be addressed are dependent upon the source of the two systems being integrated. Those sourced from a single supplier or multiple suppliers with a formal development relationship are likely to have a point integration that meets the majority of the interface requirements.

Other combinations of systems will face similar integration challenges to those outlined above.

Other Considerations

Choice of Database

A problem management system, in common with other information systems, must have some mechanism for storing the data in the system. The most common systems utilize a highly scalable relational database system to store data, and in selecting a system it is normally prudent to make sure that the system supports the more popular choices available in the market.

All systems including management systems themselves carry an administrative overhead, and aligning systems to share common backend data stores makes the day-to-day administration of the environment easier to manage.

High Availability

In designing a problem management system it is important not to lose sight of the fact that the system itself must be subject to administrative activities and may also have its availability affected by incidents in the environment.

System architecture should therefore incorporate primary and backup instances of the problem management system and proven mechanisms for failover and failback of the solution. In selecting a system, the capabilities of the system to support a high availability mode of operation should be assessed. This assessment should not simply be restricted to the main functional components of the system but should also incorporate supporting components, be they system-level components or those that control product licensing.

Utilizing Other Tools to Support Problem Management

As has been indicated in previous sections, a comprehensive problem management process must embrace both reactive and proactive postures with respect to problem management. Problem control involves the examination of past incidents to identify potential problems, while proactive problem management is more focused on trend analysis to minimize the incidence of problems in the future. Support of both of these postures requires the use of tools associated with the day-to-

[1] Gartner, 2006

day management of different process areas. The following sections attempt to explore this issue in greater depth.

Problem Control

The most commonly utilized tool for reactive problem control is the incident management system. Within the context of problem management, the incident management system is a resource that allows historical analysis of recorded incidents to identify patterns of occurrence that may indicate the presence of problem conditions that require further investigation.

The incident management system is also leveraged to promote specific classes of incidents into the status of problems. While the incident management system does serve to identify potential candidates for the problem management process, there is a wider group of tools that can assist in the investigation stage. This includes the network and element management systems associated to the configuration item affected by the problem, the change management system, capacity/performance management system, and the knowledge bases or field notices maintained by parties external to the organization such as vendors or user groups.

As the reactive problem management processes are largely focused on the examination of historical data, it is important that the supporting systems maintain a sufficiently detailed set of data that can be used to help identify a probable root cause for a particular problem. Also important is the need for all supporting systems to share a common time reference through technologies such as Network Time Protocol (NTP). Having all systems share a common time reference means that information in multiple management systems can be compared when trying to determine root cause.

Proactive Problem Management

Proactive problem management processes aim to prevent the occurrence of incidents by identification of weaknesses in the infrastructure. As such, the types of tools that support this function are those that can provide information about the quality of the services being delivered. Service-level agreement reporting tools provide a high-level view of the availability of services. Services that fail to meet their SLA targets are normally candidates for this type of proactive investigation.

While SLA reporting tools can identify the potential areas for investigation, they do not provide the granularity of information necessary to investigate issues and implement required changes. Using the problem management system to coordinate and track the investigation is the first step in managing proactive problem management effectively. The use of performance and configuration management systems is the second step in the investigation, helping to narrow down the areas of investigation to particular locations or classes of configuration items. These systems should be configured such that reporting can be used to identify trends in the operation.

The use of tools that allow stress testing of configuration items should also figure in any proactive problem management strategy. Incidents are often best avoided by making sure that proactive problem management utilizes tools that provide a framework for regression testing of any planned changes to the state of configuration items.

Meeting Problem Management Challenges

As this paper has illustrated, effective problem management requires the integration of a number of different process groups and the ability to effect change in both the infrastructure and organizational behavior. Implementing a problem management strategy that achieves these goals needs careful planning and long-term commitment. Using external supplier relationships to make

sure that the support required for both identifying problems and supporting the resolution is in place is also an integral part of the overall solution.

Cisco has long recognized that supporting problem management is critical if customers are to be successful. Resources such as the Cisco Technical Assistance Center (TAC) and other service offerings such as Network Operations Services (NOS) and Focal Technical Support (FTS) are specifically designed to provide customers with access to the expert skills that can support effective problem and error control. The frequent publishing of field notices and other classes of advisory notifications also helps customers in the proactive management area.

The rate of change or expansion of the network infrastructure and the organizational and skills issues can prevent some customers from achieving an effective problem management process. Here the need is often for externally sourced services that can bridge the gap between a current and future mode of operation. Cisco has recognized the need and has supported customers through major transitions by providing a range of professional services that can assess issues affecting the current operational state of the infrastructure and its availability and helping to define a future operational model and the resources required to support its implementation.

Implementing an effective problem management process is a challenge, but a challenge that must be met if the utility of the network as a business tool is to be fully achieved. Cisco stands ready to support all customers endeavoring to achieve higher levels of availability through operational excellence.

References

Service Operation, published by the Office of Government Commerce, Crown Copyright 2007.

Service Support, published by the Office of Government Commerce, Crown Copyright 2001.

IT Problem Management, by Gary Walker, Prentice Hall, 2001

Foundations of IT Service Management based on ITIL, published by Van Haren Publishing, 2005.

ITIL Practitioner Support and Restore Classroom Handbook, published by ITpreneurs, 2006.

ITIL is a registered trademark of Office of Government Commerce (OGC), Rosebery Court, St. Andrew's Business Park, Norwich, Norfolk, NR7 0HS, U.K.

Glossary of Terms

CCIE: Cisco Certified Internetworking Expert

CCNP: Cisco Certified Network Professional

CIO: Chief information officer

CMDB: Configuration management database

CMS: Configuration management system

CPU: Central processing unit

ITIL: Informational Technology Infrastructure Library

KEDB: Known error database

LAN: Local area network

MPLS: Multiprotocol Label Switching

MTTR: Mean Time To Restore

NOC: Network operation center

OGC: Office of Government Commerce

RCA: Root-cause analysis

RFC: Request for change

ROI: Return on investment

RSVP: Resource Reservation Protocol

SLA: Service-level agreement

WAN: Wide area network

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C11-425897-00 09/07

