Realtime publishers

"Leading the Conversation"

The Shortcut Guide™ To

Improving IT Service Support through ITIL

sponsored by

hp invent

Rebecca Herold

Copyright Statement

© 2007 Realtimepublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimepublishers.com, Inc. (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtimepublishers.com, Inc or its web site sponsors. In no event shall Realtimepublishers.com, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtimepublishers.com and the Realtimepublishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtimepublishers.com, please contact us via e-mail at info@realtimepublishers.com.

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library. All leading technology guides from Realtimepublishers can be found at http://nexus.realtimepublishers.com.]

Chapter 3: Effective Incident and Problem Management Through ITIL

Before embarking on a discussion of Incident Management and Problem Management, it is good to do some level setting. Three process terms often used within Incident and Problem Management are incidents, problems, and errors. How are these different from each other? Let's take a look at each one separately and establish parameters for each.

Incidents

The ITIL Service Support book (http://www.itil.org.uk/support.htm) defines an Incident as "Any event which is not part of the standard operation of a service and which causes, or may cause, an interruption to, or a reduction in the quality of that service.

Examples of incidents are:

- A user encounters an error when trying to access an application on the network
- Part of the WAN becomes unavailable, resulting in some users being unable to log onto the network
- Users do not get their expected messages because the email server rejects all incoming messages

Problems

The ITIL Service Support book defines a Problem as "An unknown, underlying cause of one or more incidents. A single problem may generate several incidents."

Examples of problems are:

- An application update may have made the application unusable under the same settings as before the update
- A newly installed WAN component may not be working correctly
- The ISP may not have renewed the domain name correctly

Errors

The ITIL Service Support book defines an Error as "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:

- The network settings for the desktop or server may have been misconfigured
- A network-monitoring tool may incorrectly flag a WAN circuit as being busy
- The spam filter on the email server may have been configured incorrectly

Relationship Between Incident and Problem Management

There is a close relationship between incidents, problems and errors:

- Incidents often indicate problems
- Problem investigation often leads to the identification of errors
- Errors that are unresolved can cause incidents and problems

To demonstrate this relationship, consider a common scenario within IT shops. The Service desk receives a call from an end user who got an error message when trying to log into the network. The Service desk logs the report to the incident database. An automated trend analysis determines whether this same incident has been reported, taking into consideration the time, date, and other related details about the incident. The resulting trend analysis is sent to the Problem Management system where commonalities between this and the other reported incidents can be identified. Common failures and configuration items (CIs) are identified and matched with known errors. The Problem Management system will provide a workaround or a temporary fix so that the user can get logged into the network as soon as possible. In the meantime, a request for change (RFC) may be generated to resolve the error. If the number of incidents continues to increase, the priority for implementing the RFC will become higher. When the change is implemented, the Known Errors Database will be updated to indicate the error has been resolved. Figure 3.1 shows the relationships between incidents, problems, and errors.

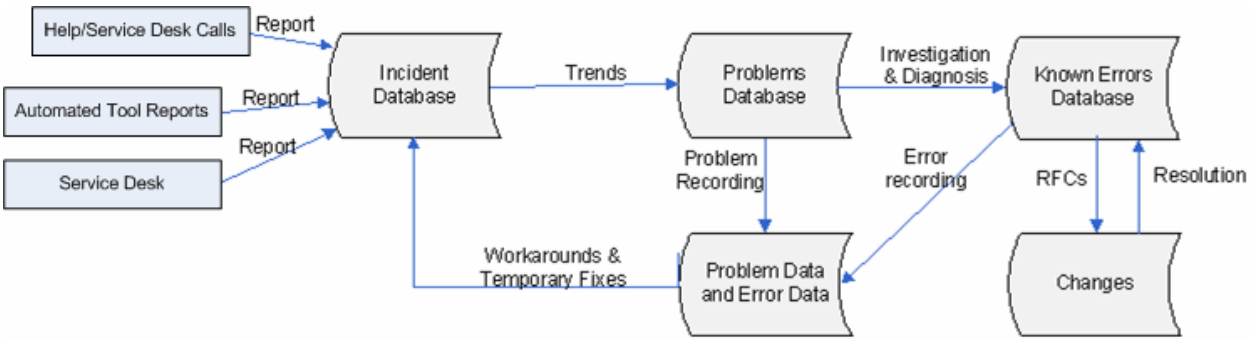

Figure 3.1: Relationships between incidents, problems, and errors.

Because of these close relationships, it is intuitive to discuss all three together.

Why Is Incident Management Important?

More incidents are reported daily. As long as technology continues to evolve, more errors will be created and incidents will continue to occur.

> Incidents impact all levels and parts of an organization. Organizations must be prepared to deal with incidents or the impact will be much more significant compared with an organization that performed no pre-planning.

Incident Management is inherently reactive. With regard to IT incidents, the goal of Incident Management is to reduce or eliminate the effects of actual or possible troubles in IT services to ensure users can get back to work, and the business can get back to being productive, as soon as possible. Incident Management has a short-term focus on restoring service.

> Information Management activities include:
>
> - Incident detection and recording
> - Classification and initial support
> - Investigation and diagnosis
> - Resolution and recovery
> - Closure
> - Incident ownership, monitoring, tracking, and communication

To most effectively address incidents, they need to be recorded and classified and the resolution for each assigned to the appropriate, qualified personnel. Incident resolution must be monitored consistently and closely to ensure incidents have been completely addressed.

The Incident Management Process

There are seven or eight steps within the Incident Management process, depending upon whether the incident involves a Service Request. Figure 3.2 demonstrates the Incident Management process.

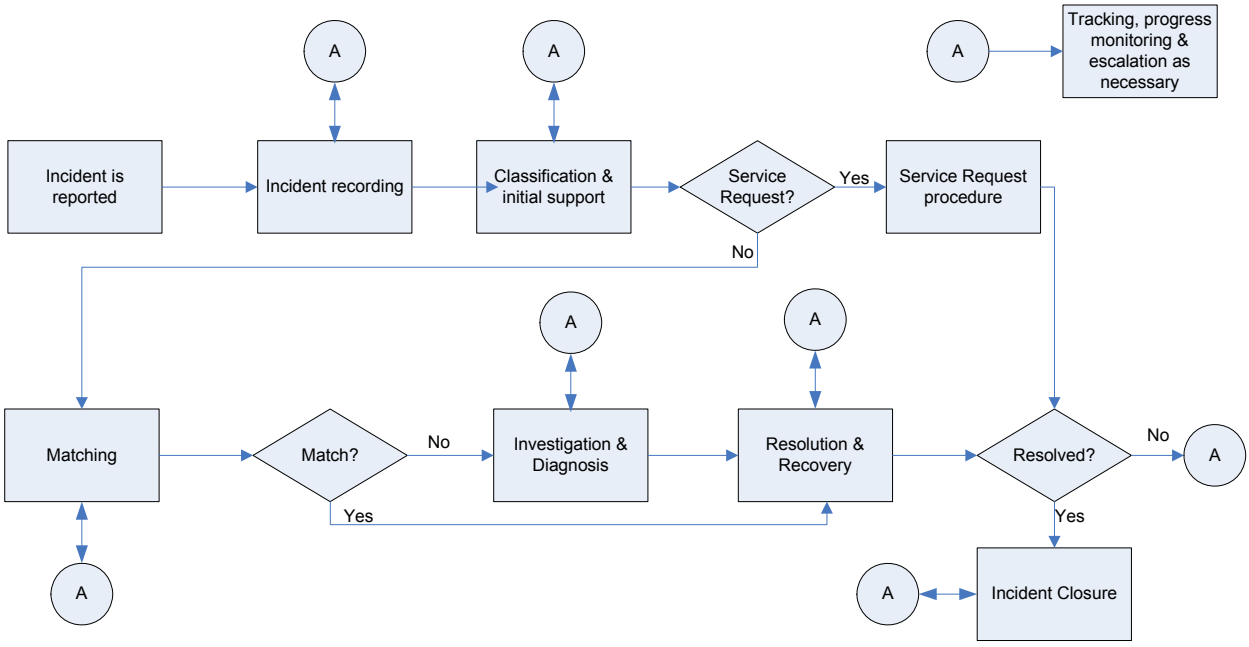

Figure 3.2: The Incident Management process.

> According to the Office of Government Commerce (OGC) Best Management Practice (http://www.best-management-practice.com/gempdf/ITIL_Glossary_V3_1_24.pdf), Service Request is defined as "A request from a User for information, or advice, or for a Standard Change or for Access to an IT Service. For example to reset a password, or to provide standard IT Services for a new User. Service Requests are usually handled by a Service Desk, and do not require an RFC to be submitted."

Incident Reporting

Incidents can be reported from any part of the enterprise as well as a number of sources outside the organization. Following a well-thought-out repeatable process will not only make incident responses more efficient, it will help to prevent similar incidents from recurring.

When the incident is reported, it is important that the details of the incident are first recorded as soon as possible. If you try to jump headfirst into incident response thinking you can always come back later and record the details, it is likely that documentation will never occur. It is also important for successful resolution of the incident that ongoing recording of significant details occurs so that progress can be accurately monitored. This documentation will also assist with addressing other incidents; learn from your experiences!

> Failure to record the incident details will not allow you to monitor compliance with SLA levels.

An important note to make about incident reporting is that each incident should not be recorded in the system more than once. Doing so will skew the incident reports and make your key performance indicator (KPI) metrics inaccurate. A KPI is a valuable metric that indicates the performance level, or success, of a particular operation or process. Management can use KPIs to make better decisions about IT processes and systems.

> 🕮 The OGC Best Management Practice (http://www.best-management-practice.com/gempdf/ITIL_Glossary_V3_1_24.pdf) defines a KPI as "A Metric that is used to help manage a Process, IT Service or Activity. Many Metrics may be measured, but only the most important of these are defined as KPIs and used to actively manage and report on the Process, IT Service or Activity. KPIs should be selected to ensure that Efficiency, Effectiveness, and Cost Effectiveness are all managed."

Classification and Initial Support

Often overlooked in typical incident response plans is classification of the incidents. Classification will allow the incident to be categorized and assist with monitoring and reporting. To create your classifications, use the following parameters:

- Category—This will include information about the origin of the incident or the support group involved. Examples include such things as processor, network, workstation, organization, procedure, Service Request, and so on.
- Priority—This will determine how quickly the incident should be addressed.
- Service—This will provide information about the services involved with the incident as covered within the associated SLA.
- Support group—This is the group that will assist with incident resolution if the Service desk cannot resolve it.
- Timelines—This will indicate the estimated time it will take to resolve the incident along with planned update times.
- Incident reference number—Assign a number not only to make it easier to find the incident data within your Incident Management system but also to reference.
- Status—Update the status to show where you are within the incident resolution process.

Matching

After the incident is classified and all associated data recorded, check to determine whether this type of incident has occurred before. If so, you can streamline the incident response time by seeing what the solution or workaround was for the previous incident and possibly use the same one, depending upon the symptoms or causal problems and/or errors.

Investigation and Diagnosis

If the Service Desk passes an incident on to a support group, the group will investigate the incident and perform diagnosis to provide resolution. If the initial group cannot resolve the incident within the targeted timeframe, they will pass it on to another support group. This will continue until the incident is resolved.

Resolution and Recovery

When the incident has been successfully solved, the support group will record all the details about the resolution into the system. If a change must occur to prevent a similar incident from recurring, a request for change (RFC) will be submitted into the Change Management process.

> It is possible that you may have an incident that does not get resolved. In this hopefully rare situation, the incident will remain open.

Incident Closure

The support group will send notice to the Service Desk that the incident has been resolved. The Service Desk will then check with the person that reported the incident and ask him or her to check the related application or system to ensure that, from their point of view and experience, the incident truly has been addressed correctly. The incident record should be updated to indicate what final category the incident is now in, along with the SLA-related metrics.

Throughout the Indicate Management process, the Service Desk is responsible for monitoring progress and updating users and customers of incident resolution status and escalation to other support groups.

Incident Management Benefits

If a well-defined Incident Management process does not exist, there is no clear accountability or responsibility for monitoring and appropriately responding to incidents. With a lack of planning and responsibility, incidents that may have been quickly resolved with a formal structure in place could become unnecessarily expansive and severely damage business by reducing service levels and leaving customers confused because they don't know what to do. In these situations, you often either have many people working on the incident—causing duplicated and often conflicting activities to occur—or you have no one addressing significant issues associated with the incident, prolonging incident resolution to an unacceptable time period. The resulting costs of the incident not only to IT but to all business customers will be much higher than it would have been if an Incident Management process were in place.

Having a well-defined, documented, and implemented Incident Management process in place not only benefits the IT areas; it benefits all areas of business. The business benefits by realizing:

- More efficient, effective, and expedient incident resolution, which reduces the negative business impact of incidents
- More productive personnel, as a result of less downtime from incidents
- Incident monitoring performed independently and is customer-focused
- SLA business management information is available
- SLA compliance

The IT area benefits will include:

- More efficient and effective use of personnel time
- Documented tracking of incidents and service requests with lessened likelihood of losing or incorrectly documenting incident information
- The CMDB is more accurate, with incident information keeping it updated as well as audited with the incident data being recorded and mapped to CIs
- The ability to improve monitoring of and measurement for meeting SLA requirements
- Better management of SLA reporting and service quality
- Customers are happier with IT services because of more effective response to incidents and less downtime

Incident Management Inputs, Outputs, and Relationships

The inputs for Incident Management are pretty clear-cut: incident data. Incidents can occur within any level or par of the enterprise. Although incidents are commonly reports by end users, the incident notification can come from a wide range of sources:

- End users
- Business leaders
- IT leaders
- External customers
- Business partners
- Automated tools

To make incident response as effective and efficient as possible, there should be a basic core of information consistently collected about each incident. These data items will determine the classification of the incident and will contribute to determining the urgency and speed for which the incident should be addressed. The data items will also support how the incident is monitored and provide information for the incident report.

Table 3.1 provides the items that should be collected when an incident is reported; these are the details that are input to the Incident Management process.

Input Item	Description
Category	Each incident should be assigned to a category and subcategory to correspond to the incident origin and support group. The following are examples of categories that can be used: • Central processing—Application, system, mainframe • Network—IP address, segment, router, hub • Organization and Procedures—Communication, order, request • Service Request—From the Service Desk • Use and Functionality—Availability, backup, capacity, service • Workstation—Keyboard, monitor, CPU, storage drive
Priority	Each incident needs to be assigned a priority to help the support groups understand which incidents need to be addressed immediately versus those that can be addressed at a later time. Priority is often computed by taking a number assigned to Urgency multiplied by a number assigned to Impact. For example, if the Urgency is 1 and the Impact is 2, the Priority is 1 × 2, or 2. Another incident may have an Urgency of 3 and an Impact of 1, so the Priority would be 3 × 1, or 3.
Service	This is a list to identify the services related to the incident. These should reference the applicable SLA requirements. Included within this list will be the escalation times for the services required by the SLA.
Support Group	If the Service Desk doesn't resolve the incident within the SLA time requirements, a support group may be called on to address the incident.
Timelines	The consideration of the SLA requirements with the priority will be used to determine the timelines for incident resolution. These need to be recorded.
Incident reference number	Each incident is assigned a reference number for easy and future reference.
Status	The status, also referenced as workflow position, indicates where progress is within the workflow. Status labels could include such terms as new, accepted, planned, assigned, active, suspended, resolved, closed, and so on.

Table 3.1: Incident Management inputs.

The escalation of an incident from the Service Desk to a support group is often described as *functional escalation*.

Outputs

Figure 3.3 illustrates the inputs and outputs for the Incident Management process.

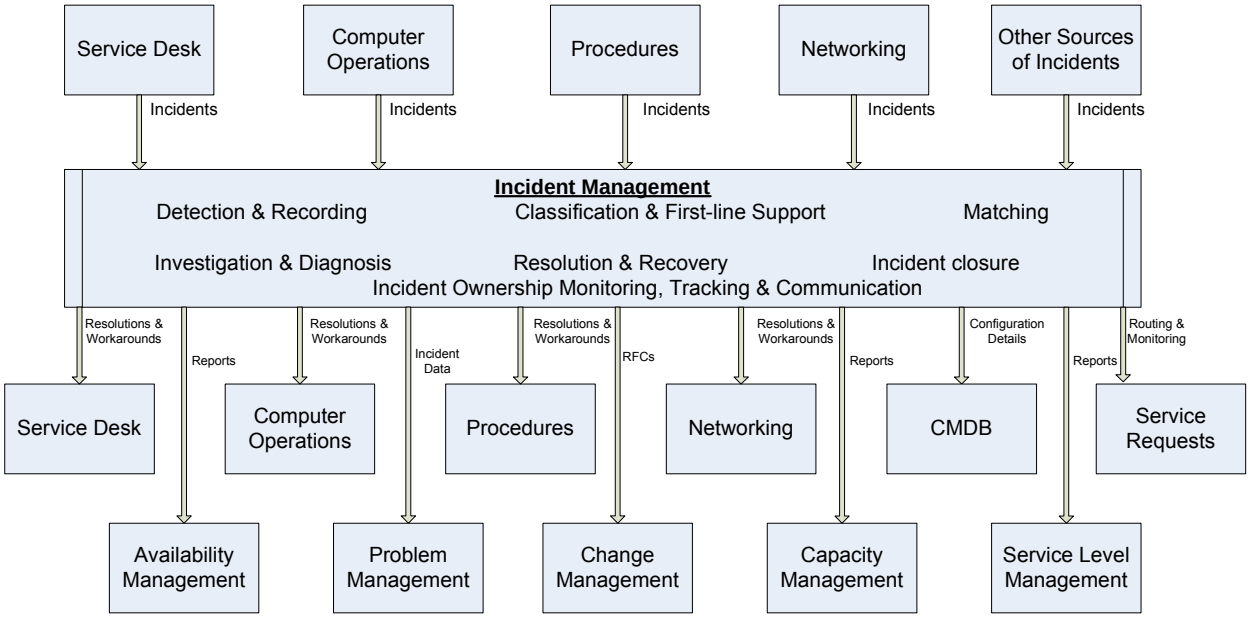

Figure 3.3: Incident Management inputs and outputs.

Relationships

Incident Management has relationships with most of the other ITIL processes. It is important for the success of not only Incident Management but of enterprise-wide ITIL processes that these relationships are appropriately managed. Figure 3.4 illustrates these relationships at a high level.

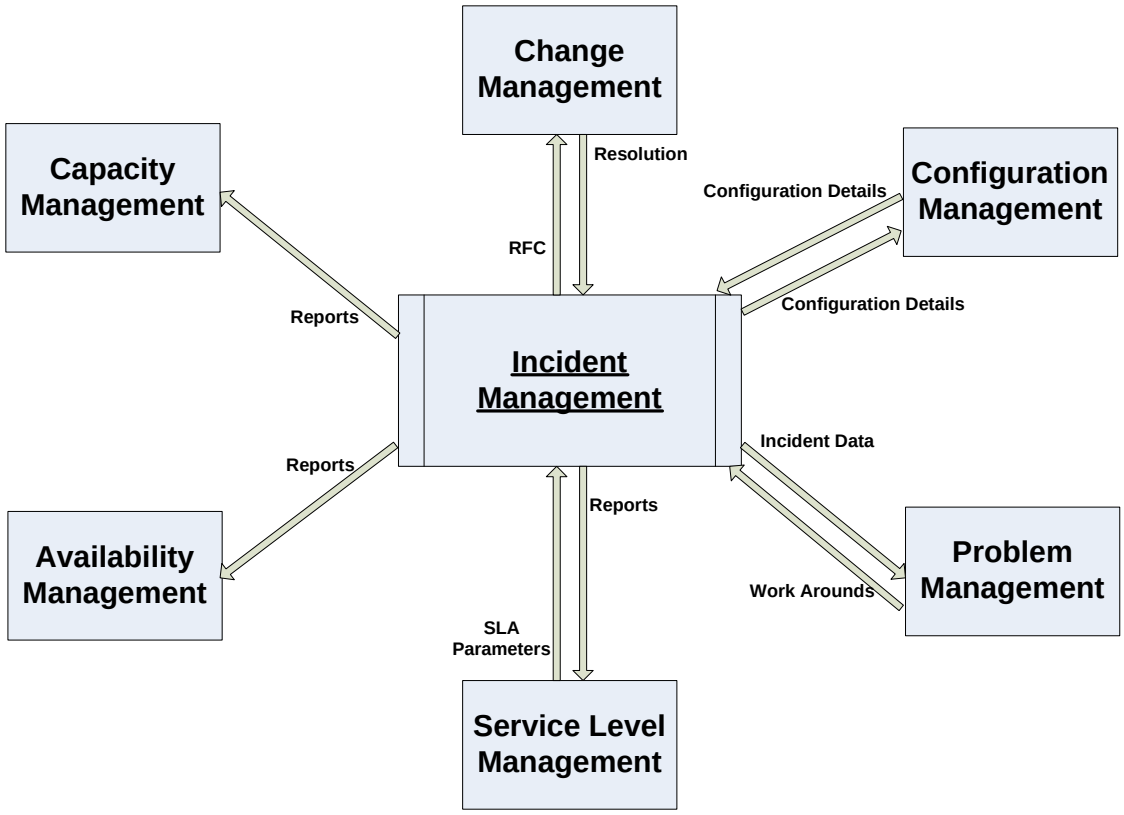

Figure 3.4: Incident Management relationships with other ITIL processes.

As this figure shows, Incident Management activities impact other ITIL processes in one way or another. It is important for effective communications channels to exist to communicate key activities. Table 3.2 provides the high-level descriptions about the relationships between Incident Management and the other ITIL processes.

ITIL Process	Relationship with Incident Management
Availability Management	Availability Management uses incident data and records in conjunction with status monitoring data from Configuration Management. Based upon the information, a service can be assigned a status, just like a CI in the CMDB. Information provided by Availability Management records can be used to determine the availability of a service and the response time of the service provider.
Capacity Management	Capacity Management uses information about incidents that are associated to capacity (for example, incidents resulting from lack of storage space, unacceptably slow response times, and so on). These events can send a notice to the Incident Management process via systems managers, business managers, or using automated tools
Configuration Management	The CMDB defines the relationships between resources, services, users, and Service Levels. Because Configuration Management defines the position responsible for each infrastructure component, incidents related to specific components can be most efficiently addressed. The CMDB can also be used to develop workarounds, such as diverting traffic to a different email server or temporarily placing a defined user group on a different print server.
Problem Management	Problem Management provides requirements for the quality of incident documentation and records that assist with determining the causal errors. It provides information about problems, known errors, temporary fixes, and workarounds.
Change Management	How are many incidents resolved? By making changes, such as replacing faulty network components or modifying parameters. Change Management provides information about scheduled changes, change status, and so on that Incident Management needs to determine appropriate actions. Additionally, changes can cause incidents. When this happens, Incident Management will send information and data to Change Management about the incidents.
Service Level Management	Service Level Management is involved with monitoring the customer agreements to ensure support provided meets customer expectations. Incident Management must understand the SLA to ensure this information is considered and used when communicating with users about incidents. Incident reports can also reveal whether service levels are provided accordingly.

Table 3.2: Incident Management relationships.

Measuring Incident Management success

Incident management metrics can help improve business. To demonstrate this, it is important to create statistics and metrics to clearly show the improvements. Success must be documented in terms of improvements to the business.

Yes, the mantra still applies; you cannot manage what you cannot measure. What kind of incident management measurements and associated data can be used to measure improvements? The following list highlights the common incident management metrics typically available for you to consider and build upon:

- Total number of incidents reported
- Total number of unique incidents
- Total number of Severity 1 incidents
- Total number of Severity 2 incidents
- Total time to resolve Severity 1 incidents
- Average time to resolve each Severity 1 incident
- Total time to resolve Severity 2 incidents
- Average time to resolve each Severity 2 incident
- Number of incidents resolved within SLA parameters
- Total number of High Severity incidents
- Total number of incidents with customer impacts
- Number of incidents reopened
- Total available non-Service Desk labor hours available to work on incidents
- Total non-Service Desk labor hours used resolving incidents
- Incident Management tools support level
- Incident Management process maturity

So where do you find this data? It can be found in such places as:

- Incident management system reports
- Labor reports
- HR reports
- Process Assessment Audit Report Findings
- Tool Assessment Results

What kind of evaluations can you make from these seemingly nondescript numbers? What are your KPIs? Some of these numbers stand on their own to provide meaningful KPIs, such as:

- Total number of incidents reported
- Total number of unique incidents
- Total number of Severity 1 incidents
- Total number of Severity 2 incidents
- Total time to resolve Severity 1 incidents
- Total time to resolve Severity 2 incidents
- Number of incidents resolved within service level agreement parameters
- Total number of High Severity incidents
- Total number of incidents with customer impacts
- Total available non-Service Desk labor hours available to work on incidents
- Total non-Service Desk labor hours used resolving incidents

However, you can do a little math and determine additional useful KPIs. The following are just some of the metrics you can calculate from the data.

Incident Resolution Efficiency Rate

You can determine the incident resolution rate by dividing the total number of incidents resolved within SLA parameters by the total number of incidents reported. For example, if there were 15 incidents reported this week and 12 of them were resolved within the SLA parameters, your resolution efficiency rate is 12/15 or 80%. This will tell your management how successful you are at resolving incidents in alignment with business requirements. The lower your efficiency rate goes, the more evidence you have that you do not have the resources or tools necessary to appropriately resolve incidents or that your SLA parameters are not realistic.

Customer Incident Impact Rate

You can determine the impact of incidents upon customers by dividing the total number of incidents with customer impact by the total number of incidents reports. For example, 15 of 20 incidents reported during the week noticeably and measurably impacted customers, such as making services unavailable, damaging business files customers depend upon, and so on, you would have a 15/20 or a 75% customer incident impact rate. This metric will tell you how successful you are at keeping incidents from impacting your customers and can point to where stronger controls are necessary, where systems need to be adjusted, and so on.

Incident Reopen Rate

You can determine the incident reopen rate by dividing the total number of incidents reopened by the total number of incidents reported. For example, if you had 5 incidents reopened during the week, and the total number of incidents reports was 20, your incident reopen rate would be 5/20 or 25%. This metric will tell you how successful you are at permanently resolving incidents. If your incident reopen rate is high, you need to look at you incident response procedures and tools and make changes to lower the rate.

Incident Labor Utilization Rate

A very useful metric to reveal how changes impact business productivity is the change incident rate. This metric will tell you how much available labor was used handling incidents. You can calculate this by taking the total labor hours (not part of the Service Desk) used to resolve incidents divided by the total available labor non-Service Desk labor hours to resolve incidents. For example, if 55 labor hours were used during the week to resolve incidents, and you had 50 hours available to work on incidents, you would have an incident labor utilized rate of 55/50 or 110%. You were over-utilized this week in working on incidents. You should keep you eye on this number to determine whether you are consistently or often over-utilized. This will help you to decide whether you should add personnel who have responsibilities for handling incidents.

Metrics such as these will tell you, and more importantly tell your business leaders, how efficient your Incident Management process components are and where improvements are needed.

Why Is Problem Management Important?

The Problem Management process includes the activities taken to minimize the adverse impacts of problems upon the business that were caused by errors within the IT infrastructure and to prevent recurrence of incidents related to these errors. Problem Management strives to get to the root cause of problems, identifies workarounds or permanent fixes, and eliminates errors.

Problem Management activities include:

- Problem control
- Error control
- Proactive problem prevention
- Providing information

Whereas Incident Management is reactive, Problem Management is primarily proactive by taking actions to determine the reasons why there was a failure in the provision of IT services. However, there are some significant reactive actions within Problem Management, such as identifying the cause of previous incidents and providing recommendations for removing those causes. Problem Management is basically an investigative process whereas Incident Management is basically a resolution process.


Many errors may be the cause of a problem. Many problems may be the result of one error.

Problem Management seeks to identify the cause or causes of a problem. The determination of the cause becomes a known error. An RFC can then be submitted to eliminate the known error along with the associated problem or problems.

The Problem Management Process

There are four basic activities involved with the Problem Management process:

- Problem control
- Error control
- Proactive problem management
- Information generation

Problem Control

Problem control activities seek to identify problems and determine the root cause of the problems. Once the causes are known, the problems can be turned into known errors that are associated with the base cause of the problem and an associated workaround. Any incident could have associated problems if the cause of the incident is not known.

The first step in problem control is identifying the existence of a problem along with recording significant details about the problem. The problem should then be classified according to the

- Appropriate category, such as hardware or software
- Impact upon the business and associated business process and applications
- Priority based upon consideration of urgency, impact, risk, and the sources necessary to resolve the problem
- Status of the problem
- Urgency of finding a solution

The classification of a problem may change throughout the process of resolving the problem. For example, implementing a temporary fix or using a workaround may lessen the urgency and impact.

In addition to classification, an impact analysis should be performed to determine how serious the problem is and what potential and actual effects the problem has on IT services. This impact analysis will become the basis to mitigate and manage the risk. Based upon the results of the impact analysis, a priority is assigned to the problem and then the appropriate personnel and resources can be assigned to resolve the problem.

The problem is now ready to be investigated and diagnosed. Investigation and diagnosis will typically need to be repeated multiple times. Each time you will get closer to resolution. Too many IT practitioners believe that resolution should or can occur quickly, but with this attitude, you will be setting yourself up for failure and frustration.

Investigation often includes trying to reproduce the problem within an isolated environment. This is a very good tactic. Don't be afraid to call in specialists from the support group to help.

If an acceptable workaround can be established after the cause of the problem is discovered, and the CIs responsible are identified, a relationship between the incident and CIs will allow for a known error to be defined. If an RFC must be submitted to apply a temporary fix, the RFC process must be followed.

Error Control

The error control activities involve monitoring and managing all known errors from the time they are identified until they are resolved. Many areas throughout the enterprise may be involved with error control.

When the cause of the problem has been determined and the corresponding CIs identified, the problem can be linked to a known error, which launches the error control process. At this point, data is sent to the Incident Management process to use within any open incidents. An existing workaround for the known error can also be used to assist with incident resolution.

The team working within the Problem Management process will determine what needs to be done to resolve the problem if the errors are known. The team members should compare the possible solutions and choose the one that is the best fit with the associated SLAs, costs, impacts, and urgency. When the decision has been made regarding the best solution for resolving the problem, an RFC can be submitted to Change Management.

Although most problem and failures are identified in the production environment, it is important to keep in mind that test and development environments can also have failures and known errors.

When the changes to fix the error have been implemented, a Post Implementation Review (PIR) should be done before closing the problem. Incident Management should be sent the results of the PIR so that they can close the applicable incidents.

Throughout the error control activities, there is constant tracking and monitoring to stay abreast of problem and error resolution. Tracking and monitoring will help determine whether the business impact and/or urgency changes, if the priority changes, and whether the RFC has been successfully implemented and addresses the problem or error.

Proactive Problem Management

The actions that occur within proactive problem management, which basically means actions taken to prevent problems, ensures the quality of the services and underlying infrastructure. Trend analysis occurs along with actions to identify weaknesses. Proactive problem management can have a huge impact on the business by identifying, investigating, and addressing weaknesses throughout the infrastructure components before they result in incidents.

Information Generation

Throughout Problem Management processes, information is generated and shared. The closest relationship is with Incident Management, to which information is passed concerning workarounds and temporary fixes. Information is also obtained from the CMDB to determine the other entities that need to receive information about the problem resolution. The SLA is also used to see what additional entities need to receive information. Figure 3.5 demonstrates the Problem Management process.

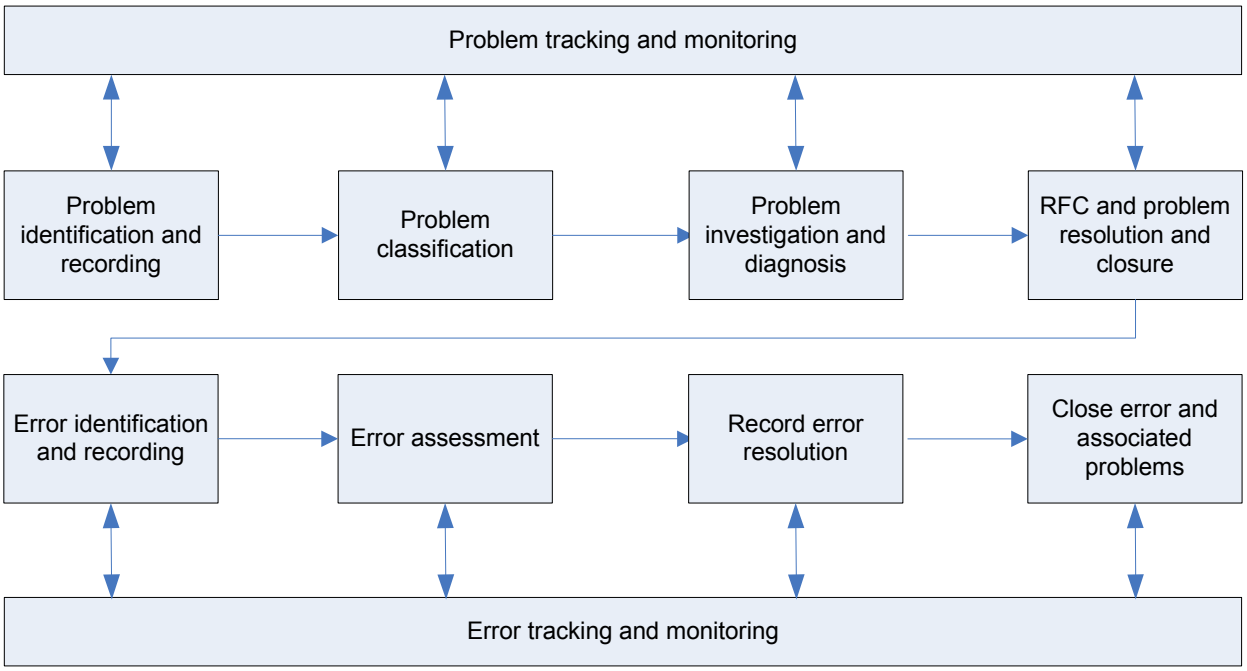

Figure 3.5: The Problem Management process.

Problem Management Benefits

The objective of Problem Management is to identify and eliminate the causes of incidents so that actions can be taken to prevent them from happening again. That alone would seem to be a compelling benefit for implementing Problem Management processes. However, in case this does not sway you, the following list highlights a few more benefits:

- Improves the quality of IT services by taking actions to reduce the number of incidents and thus reduce the IT workload.
- Improves the documentation related to problems, errors, and incidents.
- Improves user productivity by addressing and removing errors and problems, which results in giving users more time to actively perform business-related activities.
- Documentation will improve support team productivity by having documentation to reference to resolve incidents on an ongoing basis more efficiently, economically, and quickly.
- Raises the stature of IT services reputation. When IT services become more stable and systems availability increases, customers will be more willing to entrust business activities to IT areas.
- Documentation can be used to perform trend analysis that can result in implementing procedures and tools to prevent incidents. Documentation is also available and useful for investigations and in preparing RFCs.
- Establishes a standard for consistent and thorough incident and problem recording, classification, and reporting.
- Provides details on workarounds and temporary fixes allowing first-line support personnel to be more likely to resolve incidents.

Inputs, Outputs, and Relationships

Six other ITIL processes provide input to the Problem Management process:

- Incident Management provides incident record data used by Problem Management to identify problems.
- Change Management provides PIR results about associated incidents, problems, and errors.
- Configuration Management provides information critical for resolving problems, such as infrastructure details, software and hardware configurations, services, architecture blueprints, and so on.
- Availability Management provides availability design, planning, and monitoring data.
- Capacity Management provides data about storage, bandwidth settings, and other details useful for problem shooting.
- Service Level Management provides SLA data along with other quality data.

Outputs

Problem Management provides output to two other ITIL processes:

- Change Management receives RFCs to help resolve problems.
- Incident Management receives matching information to determine whether a problem has been associated with other incidents.

Figure 3.6 illustrates the inputs and outputs for the Problem Management process.

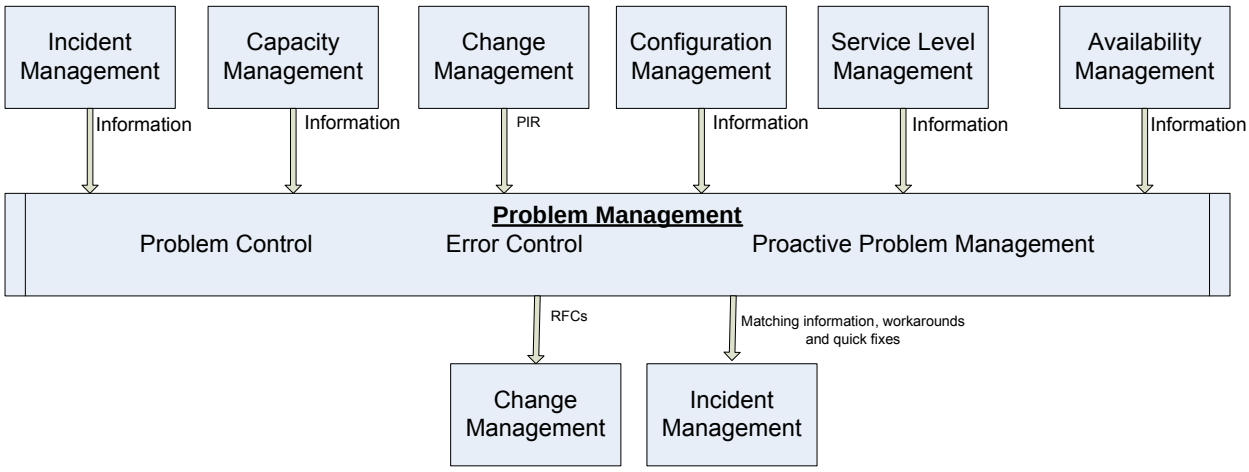

Figure 3.6: Problem Management inputs and outputs

Relationships

Problem Management has relationships with six other ITIL processes. It is important for the success of not only Problem Management but of enterprise-wide ITIL processes that these relationships are appropriately managed. Figure 3.7 illustrates these relationships at a high level.

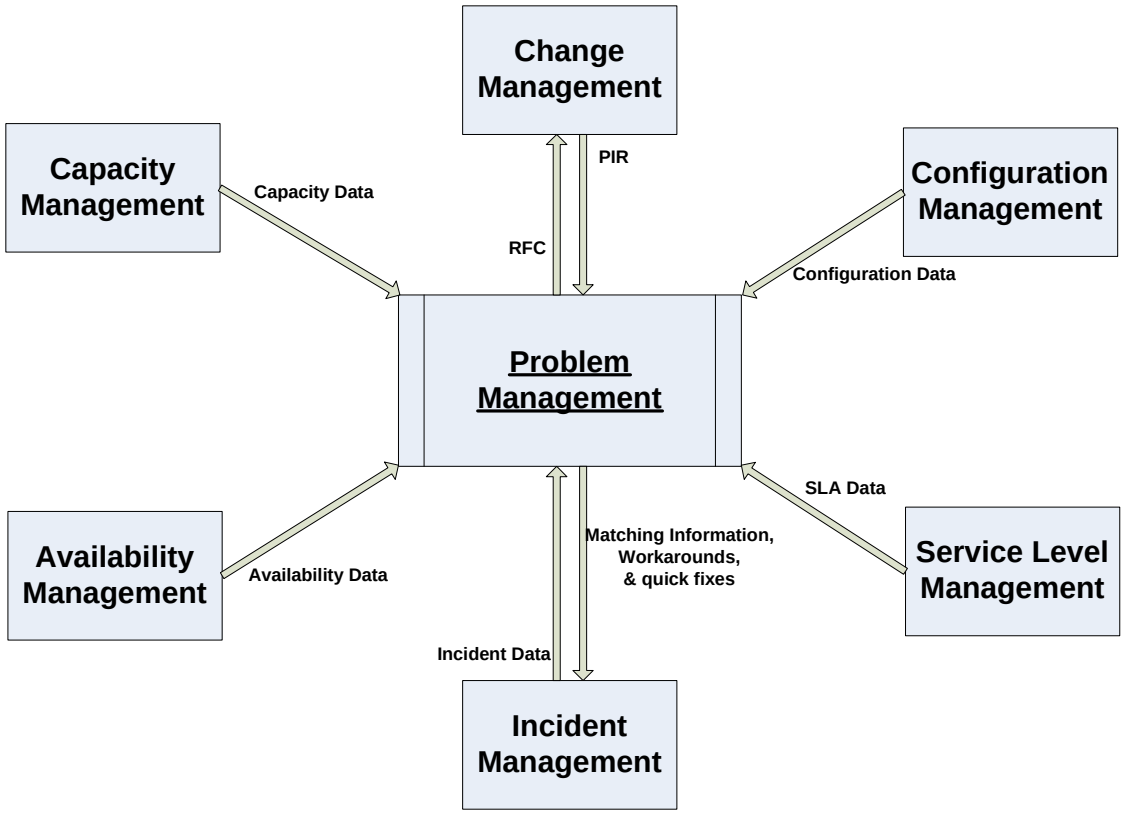

Figure 3.7: Problem Management relationships with other ITIL processes.

Table 3.3 provides high-level descriptions about the relationships between Problem Management and the other ITIL processes.

ITIL Process	Relationship with Problem Management
Incident Management	Incident Management provides incident record data used by Problem Management to identify problems. Incident Management receives matching information to determine whether this problem has been associated with other incidents.
Change Management	Change Management provides PIR results about associated incidents, problems, and errors. Change Management receives RFCs to help resolve problems.
Configuration Management	Configuration Management provides information critical for resolving problems, such as infrastructure details, software and hardware configurations, services, architecture blueprints, and so on.
Availability Management	Availability Management provides availability design, planning, and monitoring data.
Capacity Management	Capacity Management provides data about storage, bandwidth settings, and other details useful for problem shooting.
Service Level Management	Service Level Management provides SLA data along with other quality data.

Table 3.3: Problem Management relationships.

Putting Incident Management and Problem Management into Action

Let's revisit our ACME Super Duper Supplies business from previous chapters and step through an example to see how all these Incident Management and Problem Management processes are related. ACME Super Duper Supplies recently implemented a new ecommerce Web site that allows for online merchandise ordering and payments for their new product, Magic Mover. This was a significant change in their IT infrastructure in addition to having a major impact on their business. Much money was invested in this change, so there are great expectations for a large financial return.

The new product was popular right away, and the new Web site continued to get increasingly more hits from day to day. This was great news to the business unit manufacturing the product. However, as sales were ramping up to a very profitable level, the Web site crashed and Web site customers received only error messages when trying to get to the site.

The Service Desk received an automated notice that the Web site was down. They contacted Ms. Flint, the manager of the Magic Mover business unit, and notified her of the incident. They also performed first-line support to determine whether they could resolve the incident. They were not able to get the site back up within their goal timeline, so they passed the incident on to the second-line support group.

The second-line support group performed a change impact analysis and were able to apply a temporary fix and get the Web site back up and available for business. The support group enters the incident data into the Incident Management system, from which data is sent to the Problem Management system. The Problem Management team receives the information and investigates by performing deep change impact analysis to identify the root cause of the problem. They found that a known error with capacity settings triggered the Web site outage.

The Problem Management team submits an RFC to modify the capacity settings to eliminate the error and prevent the incident from recurring. A PIR is performed to determine whether the changes truly resolved the problem. Figure 3.8 shows how the Incident Management and Problem Management processes flow to address the Magic Mover Web site incident.

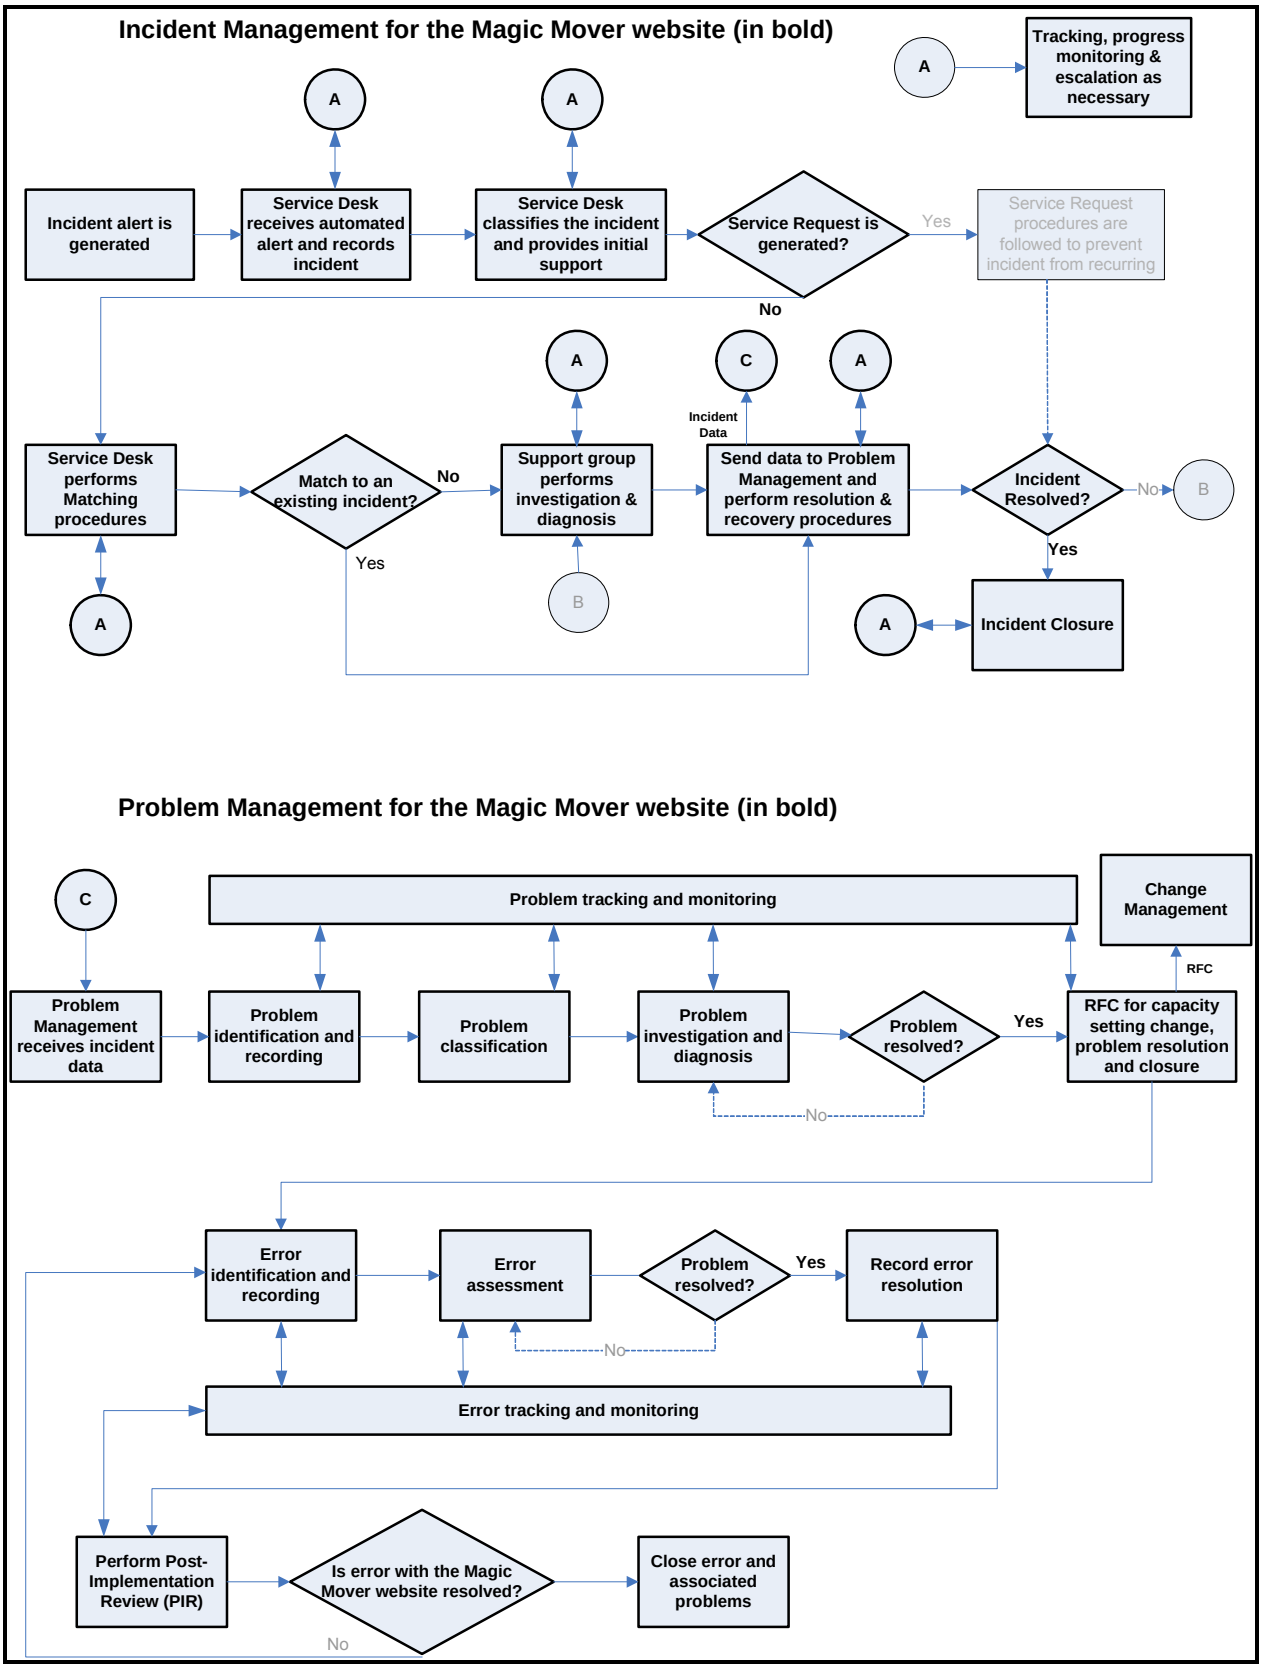

Figure 3.8: Magic Mover Incident Management and Problem Management process flow.

Costs

It is important for you to consider the costs involved with implementing ITIL Incident Management and Problem Management processes. These costs will generally fall into two categories: people costs and technology costs.

People Costs

You likely already have personnel throughout the enterprise performing Incident Management and Problem Management tasks, but in an ad hoc or otherwise uncoordinated way. If you are not already using ITIL, it is likely that they are performing these tasks, but in silos, meaning they are repeating tasks, leaving out important tasks, or performing conflicting tasks. When implementing Incident Management and Problem Management processes, you should be able to use some of these same personnel that are now freed up for implementation.

Personnel costs will include such things as the time of the personnel who are members of the support groups when they are actively resolving incidents as well as any training they need to receive. There are also personnel costs in maintaining and upgrading the associated Information Management and Problem Management systems and tools. A typically significant cost is the up-front time necessary to plan, define, communicate, and implement the Incident Management and Problem Management processes.

Technology Costs

Technology costs will include such things as tools to support the Incident Management and Problem Management processes, possibly hiring outside consultants or technicians to assist in implementation of the tools, storage space for incident data, and any training costs that may be necessary.

You will need to plan carefully the hardware and software tools you decide to use for implementing the automated portion of the Incident Management and Problem Management processes, and ensure that they integrate with the other ITIL processes. A good, integrated technology tool may be a significant up-front investment, but if chosen and implemented correctly, it will result in long-term savings in other areas of the enterprise.

Measuring Problem Management Success

Problem management metrics can help improve business. But to demonstrate this, it is important to create statistics and metrics to clearly show the improvements. Success must be documented in terms of improvements to the business.

Shall we repeat the mantra again? You cannot manage what you cannot measure. What kind of Problem Management measurements and associated data can be used to measure improvements? The following are some of the common Problem Management metrics typically available for you to consider and build upon:

- Total number of incidents reported
- Total number of incidents reopened
- Total number of major problems
- Total number of problems in the pipeline
- Total number of problems resolved and removed
- Total number of known errors
- Total number of problems reopened
- Total number of problems with customer impact
- Average time to resolve each problem
- Average time to resolve Severity 1 problems
- Average time to resolve Severity 2 problems
- Total available labor hours allotted to work on problems
- Total labor hours spent working on problems
- Problem Management tools support level
- Problem Management process maturity

So where do you find this data? They can be found in such places as:

- Incident Management system reports
- Problem Management system reports
- Labor reports
- HR reports
- Audit reports
- Process Assessment Audit Report Findings
- Tool Assessment Results

What kind of evaluations can you make from these seemingly nondescript numbers? What are your KPIs? Some of these numbers stand on their own to provide meaningful KPIs:

- Total number of major problems
- Total number of problems in the pipeline
- Total number of problems resolved and removed
- Total number of known errors
- Total number of problems reopened
- Total number of problems with customer impact
- Total available labor hours allotted to work on problems
- Total labor hours spent working on problems

However, you can do a little math and determine additional useful KPIs. The following are just some of the metrics you can calculate from the data.

Customer Impact Rate

You can determine the customer impact rate by dividing the total number of problems with customer impact by the total number of problems in the pipeline. For example, if you had 50 problems in the pipeline this week and 22 of them impacted customers, your customer impact rate is 22/50 or 44%. This will tell your management how well you are at keeping problems from impacting your customers and point to where you need more resources, tools, or labor to lower the rate to an acceptable level.

Incident Repeat Rate

When incidents repeat and must be reopened, it points to underlying problems that must be discovered. You can determine the incident repeat rate by dividing the total number of repeat incidents by the total number of incidents. For example, if you had 50 incidents during the week, and 25 of them were repeat incidents, your incident repeat rate is 25/50 or 50%. This will tell your management how effective you are at minimizing repeat incidents. The higher the number, the more investigation and research that needs to be done to determine any existing problems at the core of the incidents.

Problem Labor Utilization Rate

You can determine the problem labor utilization rate by dividing the total labor hour spent working on problems by the total labor hours available to work on problems. For example, if you spent 80 hours resolving problems during the week and you had allotted 120 hours to be available for problem resolution, your problem labor utilization rate would be 80/120 or 67%. This metric will indicate how much available labor capacity was used handling problems, and can indicate whether the number allotted is too low, if more personnel is needed, or if there needs to be changes in the procedures to fix problems.

Problem Reopen Rate

The problem reopen rate is found by dividing the total number problems reopened by the total number of problems in the pipeline. For example, if you had 60 problems in the pipeline for the week, and 20 of the problems were reopened, your problem reopen rate would be 20/60 or 33%. This metric will tell your management how successful you are at permanently removing problems. Your goal will be to get this rate as low as possible.

Problem Resolution Rate

The problem resolution rate is computed by dividing the total number of problems resolved by the total number of problems in the pipeline. For example, if you had 45 problems in the pipeline for the week, and you resolved 30 of them, your problem resolution rate would be 30/45 or 67%. This metric will tell your management the percentage of problems you successfully addressed and removed. The higher the percentage, the better.

Problem Workaround Rate

The problems workaround rate is found by dividing the total number of known errors by the total number of repeat incidents. For example, if the total number of known errors is 100 and the total number of repeat incidents is 120, your problem workaround rate is 100/120 or 83%. This metric will tell your management the percentage of problems for which you implemented workarounds.

Metrics such as these will tell you, and more importantly tell your business leaders, how efficient you are at implementing Problem Management process components and where improvements are needed.

Summary

Implementing the ITIL Incident Management and Problem Management processes will be an evolutionary process just as Change Management implementation was. It will take time and investment up front. It will be a learning experience. But, when done correctly, it will make your business more efficient; reduce downtime; prevent incidents from happening; save money that otherwise would have been spent constantly addressing recurring incidents, problems, and errors; and make IT more strategic in the eyes of your business leaders.

Incident Management and Problem Management implementation success will also take the strong, consistent commitment of your executive management to get through the inevitable growing pains. Be sure you have that to get the subsequent commitment of your ITIL team members, and ultimately improve your Incident Management and Problem Management processes.

Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit http://nexus.realtimepublishers.com.