DITY™ Weekly Newsletter
itSM Solutions® IT Experience. Practical Solutions
The workable, practical guide to Do IT Yourself Vol. 5.7 • February 18, 2009

Expanding the Expanded Incident Lifecycle

By Janet Kuhn, Vice President, itSM Solutions LLC

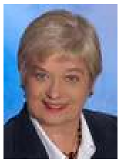

A key to improving the quality of IT service begins with understanding and utilizing one of ITIL's simplest concepts - the Expanded Incident Lifecycle.

If you have attended an ITIL Foundation course, you undoubtedly remember the slide depicting the Expanded Incident Lifecycle (Figure 1, below). That is the graphical timeline that starts with an Incident on the left, progresses through the various stages of diagnosis, repair, restoration and closure, and then continues to the next Incident.

The labels dispersed along the Incident timeline are not just handy monikers that the Service Desk uses to report the changing state of an Incident. They represent critical intersections of ITIL processes and activities and provide a roadmap to shorten the time to recover from an Incident and lengthen the time of error-free operation.

Mean Times to . . .

Before we start, let's review a few key ITIL measurements, the "Mean Times to . . ."

MTTR (Mean Time to Repair) - This is the average elapsed time between detecting an Incident and repairing the failed component; e.g., diagnosing and replacing a failed disk. Upon the completion of this activity, there is a functioning disk, but data has not been restored, and the users are still unable to access or use the service.

Essentially this measures the technical response to diagnose and repair the failed component. The shorter this time, the better because shortened times mean less downtime for the user.

MTRS (Mean Time to Restore Service) - This is the average elapsed time between detecting an Incident and fully restoring the service to the user; e.g., restoring data to the disk, recovering and restarting interfaces to other applications, informing the users that the service is available, and initiating user access (you may not want all of your users to log in simultaneously upon repair of the service!).

This is a measure of the quality of your operational processes, as well as system design to facilitate recovery after failure. Again, shortening these times should be your goal.

MTBF (Mean Time Between Failures) - This is the average elapsed time between restoration of service following an Incident and detection of the next Incident. In this case, a big number representing a long time between failures is good because it indicates a reliable service.

MTBSI (Mean Time Between System Incidents) - This is the average elapsed time between Incidents, including downtime represented by the MTTR and MTRS measurements. By understanding the proportion of repair and restoration time versus failure-free time for a particular service, you can begin to prioritize service and system improvements. For example, you may decide to commit resources to improving a critical business service that experiences few, but lengthy, failures, and give a lower priority to repairing a less business-critical service that experiences frequent failures, but which require few resources and time to restore.

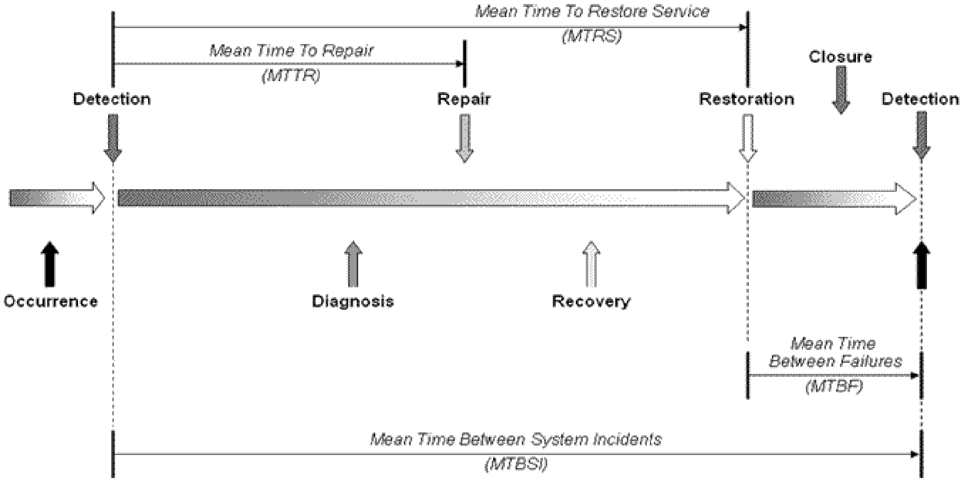

Figure 1

Expanding the Expanded Incident Lifecycle

Now that we've looked at what the Expanded Incident Lifecycle diagram tells us, let's take a look at which ITIL processes support it, and how you can use it to pinpoint areas to automate or improve.

Occurrence – By definition, an Incident is an unplanned disruption to an agreed service. ITIL offers a number of proactive ways to protect a service:

- Capacity Management – Capacity Management seeks to ensure that proactive measures to improve performance of services are implemented when it is cost-justifiable to do so.
- Availability Management – Availability Management identifies Vital Business Functions (VBF) that are critical to the business and implements designs that reduce the likelihood of unavailability.
- IT Service Continuity Management – IT Service Continuity Management (ITSCM) evaluates risks and threats to IT services and seeks to avoid them or to moderate them if they do occur.
- Information Security Management – Information Security Management proactively improves security controls, security risk management and reduction of security risks.
- Access Management – Access Management executes policies set by Information Security Management.
- Service Level Management – Service Level Management does not really prevent Incidents, but it works with the business to define the levels of agreed services.

Detection - Incident resolution starts when a user or an automated system detects an error with a Configuration Item. Detection generally occurs sometime after the occurrence of the event. The goal is to shorten the time between Occurrence and Detection as much as possible. This activity ties directly to:

- Capacity Management – Capacity Management ensures capacity can be monitored and measured.
- Event Management – Event Management establishes threshold monitoring activities to detect Incidents early.
- Incident Management – Incident Management should interface with the Service Desk and Event Management (as well as the Operations Management and Technical Management functions) to bring all Incidents under the control of Incident Management.
- Service Desk – The Service Desk should have many channels for users to report Incidents when they occur.

Diagnosis – During this stage, staff members try to identify the characteristics of the Incident and match it to previous Incidents, Problems and Known Errors. If Incident Management cannot match the Incident, the Problem Management process should start.

- Incident Management – Establish a good interface with Problem Management to match the Incident to existing Problems and Known Errors or to report a new Problem.
- Problem Management – Establish strong Problem Management procedures to rapidly and accurately diagnose problems.

- Supplier Management – Establish Supplier Management procedures that document how third-party suppliers will be involved in diagnosis activities.
- Service Level Management – Agree and establish Operational Level Agreements (OLA) with the Operations and Technical Management functions so everyone knows how to prioritize an Incident or Problem.
- Technical Management Function – Document working procedures so that all staff knows what their roles and responsibilities for diagnosing Problems are.

Repair – Sometimes a repair might raise a Request for Change (RFC) to change one or more Configuration Items (CI). After the CI is repaired, it may still be unavailable to the user and require recovery.

- Change Management – Establish strong Change Management procedures to control changes made as a result of a problem diagnosis.
- Supplier Management – Establish Supplier Management procedures for repairs that are made by third-party service providers.
- Technical Management Function – Ensure Technical Management staff have the proper levels of skills and training.

Recovery – This is the process of restoring the failed CI to the last recoverable state. This includes any required testing, final adjustment, configuration, etc.

- Change Management – Ensure Change Management includes recovery steps in its planning.
- Operations/Technical Management Function – Ensure that the Operations/Technical Management functions of Service Operation document and understand the steps to recovery.

Recovery also has a proactive side, which results in designing services and systems that faster and easier to recover.

- Problem Management – Problem Management should review and document problems and potential problems to develop proactive Problem solutions that can be shared by all IT Service Lifecycle phases and processes.
- Service Design – Establish strong Service Design processes to design services to expedite their recovery from failure.

Restoration – Service restoration makes the recovered service available to the user, so that the user can resume work.

- Service Desk – Establish a strong interface with the Service Desk to manage user communications during implementation of the change and restoration of the service.
- Service Transition – Establish a strong interface with Service Transition to implement changes and restore service to the users.
- Operations Management Function – Establish good documented procedures with the Operations Management Function to implement changes and restore service to the users.
- Technical Management Function – Establish good documented procedures with the Technical Management Function to implement changes and restore service to the users.

On the proactive side, restoration capabilities can be "designed into" the service:

- Service Design – Service Design should include restoration considerations in its analysis and design of new services.

Closure – Closure occurs some time after restoration. It should give the user ample time to "shake out" the repaired service to ensure that it is really working, but it should not be so far into the future that users and staff have difficulty reconstructing what the parameters of the actual failure were.

- Service Desk/Incident Management – The Service Desk and Incident Management process should formally close each Incident after verifying its closure with the user.
- Change Management – Change Management includes an immediate technical review to ensure that the Change has been implemented properly and does not create other problems. Later it does a long-term review to determine whether the change has created the benefits (beyond resolving a Problem) that the user had anticipated.
- Service Level Management – The Service Level Management process should agree with the business what constitutes "closure" of an Incident. Also, it includes Incidents in its periodic performance reviews with the business.

The Final Step – Closing the Loop

You do not see the “final” step in the Expanded Incident Lifecycle because it is not really a step, but an action “implied” by the four **“Mean Time”** measurements.

The last step that ties together the steps along the timeline is to use MTTR, MTRS, MTBF and MTBSI to **measure** and **analyze** the effectiveness and efficiency of all of the activities and processes that contribute to Incident restoration.

Were appropriate resources available to assist with the Incident resolution? Were appropriate interfaces in place so that resources could be applied in a timely manner?

And, finally, did you learn something that can help the process work better the next time – or prevent the Incident from occurring?

Entire Contents © 2009 itSM Solutions® LLC. All Rights Reserved.
ITIL ® and IT Infrastructure Library ® are Registered Trade Marks of the Office of Government Commerce and is used here by itSM Solutions LLC
under license from and with the permission of OGC (Trade Mark License No. 0002).