INCIDENT MANAGEMENT: THE COMPLETE GUIDE

Technology is everywhere, and we depend on it. Whether it be home, work, school, health, or civic needs, technology is involved.

That's why we get so frustrated with any disruption from normal operation. Social media is awash with stories of systems not performing as they should: banking systems, healthcare portals, airline booking, online shopping, even the social media platforms themselves.

For this reason, most relationships between any service provider—you—and your customers depend heavily on whether you can ensure minimal disruption. When the inevitable disruption does occur, you must manage the incident in a way that the consumer has agreed to tolerate.

Incident management is the formal name of this necessary business practice, and it's not one for companies to take lightly, no matter your industry. This article will look at many parts of the incident management practice, including:

- Incidents
- The practice itself
- In-depth examples
- Workflow & activities
- Best practices
- Additional resources

Let's get started.

What is an incident?

When it comes to ensuring that operational services provide value to customers, incident management is among the most important disciplines. ITIL® 4 defines an incident as:

> *An unplanned interruption to a service or reduction in the quality of a service.*

Here are some other definitions:

- ISO 20000 defines an incident as an unplanned interruption to a service, a reduction in the quality of a service, or an event that has not yet impacted the service to the customer or user.
- VeriSM broadens the term "issue" to covering situations when a customer perceives a service interruption, as well as actual interruptions.

How a service provider handles incidents plays a very significant role in determining customer satisfaction. Here are some examples of an incident in an online system:

- Users not being able to log in
- The system's lack of responsiveness to commands
- Perceived slowness compared to normal
- Corrupted or hacked data

Of course, not all incidents are visible to the end user. But they still require your attention.

What is incident management?

The purpose of the incident management practice is to minimize the negative impact of incidents by restoring normal service operation as quickly as possible.

Whether it's a crashed laptop, corrupted data or a painfully slow application, how we respond and deal with the interruption to service indicates whether we have an optimal incident management process.

This practice can be handled by an individual, teams or multiple organizations depending on the scale. Successful organizations designate a specific Incident Commander (IC)—one person who is responsible for leading a temporary cross-functional team to focus all energies and attention towards a swift resolution.

Your company's ability to quickly address incidents is a key factor in:

- User and customer satisfaction
- Your credibility and reputation
- The value you create in your relationships

Clearly, these are areas you must excel at—making incident management a critical activity. At its most essential, incident management involves two main activities:

1. Record
2. Manage

Once you identify, or get notified of, the incident, you would capture just enough information about it, including description, time, and source. This record then becomes the basis for analysis and decisions on managing the incident, including:

- Communication
- Resolution
- Escalation
- Handoff to other processes

The classic incident management activities as outlined in ITIL 4 are seen below:

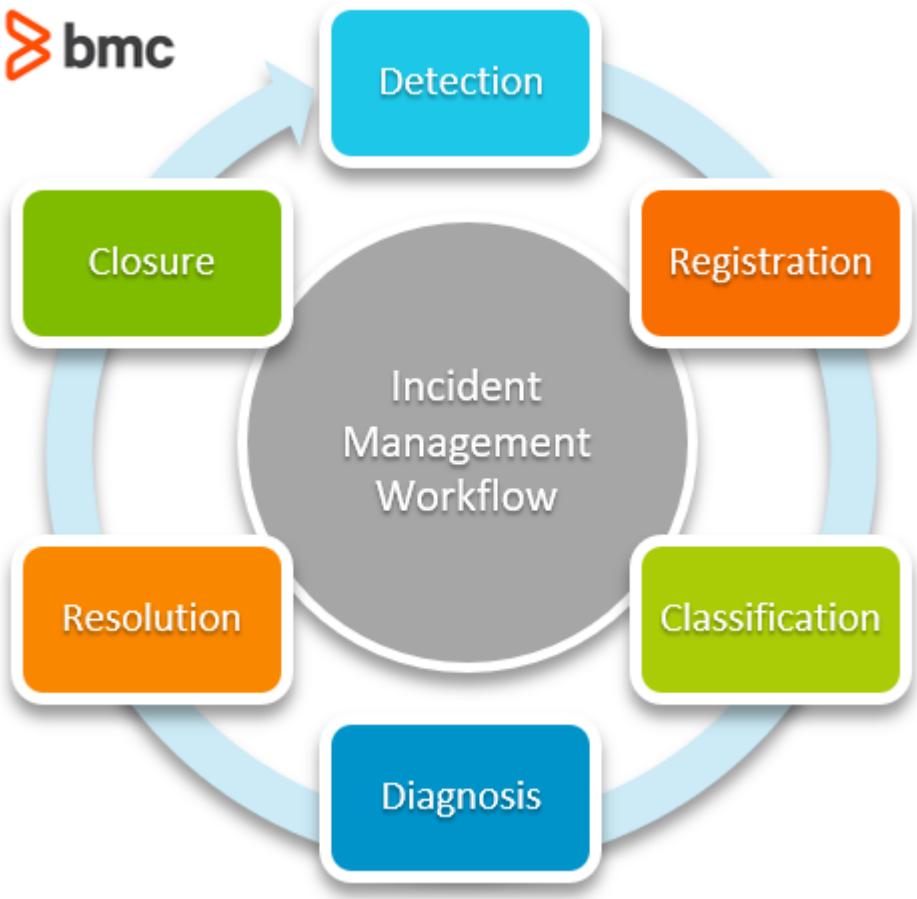

Successful incident management relies on having a clear understanding of what the customer agreed to or is willing to tolerate regarding the duration and handling of any particular incident. This is usually defined in service level agreements (SLAs) or contracts, which include timelines for responding and resolving incidents based on some criteria, usually priority, as a function of impact and urgency.

As the service provider, how you structure your organization to handle different types of incidents is a major driver in your incident management execution:

- **Some incidents may be repeatable, with their causes known.** In these cases, you can define and use incident models for handling and resolution. An incident model is a repeatable approach to managing a particular type of incident. Models help reduce both resolution time and the learning curve for new employees.
- **Where a solution to an incident is not easily found, a workaround may be applied** to try and lessen the impact and/or probability of recurrence. Workarounds, such as restarts or reconfigurations, can quickly restore services back to an acceptable level of quality.

(Importantly, incident management is different from problem management, which focuses on how you handle the problem in the future.)

Examples of incident management

We'll put this theory into practice.

These three incident scenarios can provide a good picture on how best to handle common service interruptions, using good practices and standards.

Single user-related incident

Ben, who's recently been promoted from customer care to marketing agent, is trying to log into the company's CRM to check on some new leads. Unfortunately, his credentials are not going through. He tries to reset his password and still nothing works. He decides to reach out to the IT helpdesk via phone.

Tiffany, a help desk agent, takes Ben's details and logs them into the helpdesk system, verifying that he works for the company. She then proceeds to log into the CRM administrator module and checks on his profile. It seems that some changes to his profile were not executed correctly, leading to the error.

She verifies the requested changes and proceeds to correct them. She then asks Ben to try log in, and he is successful!

Ben continues with his work as Tiffany proceeds to close the record on the helpdesk system which sends a satisfaction survey to Ben. He happily gives her a 5-star rating.

Tiffany proceeds to check a sample of related changes for CRM profiles assigned to the same agent who handled Ben's profile. Fortunately, it seems the rest were well executed. "No problem ticket required!" she happily sighs.

Multi-user service incident

The service desk manager Hilda notices an uptick in calls—her entire team is now fully engaged talking about the same thing. "It's the invoicing system. Employees are unable to submit their timesheets," remarks one of her staff.

Being the second Friday of the month, Hilda knew that the majority of staff will submit their timesheets today, so it's going to get crazy. She immediately rings up the IT manager who confirms that the system experienced a database error which they are working on. An incident ticket has already been logged on the ITSM system by his lead systems engineer.

Hilda notifies her team and then logs into the ITSM system to post a bulletin about the invoicing system issue. Her team immediately works to relate all the received records to the single incident ticket so that closure can be managed centrally. 20 minutes later, Hilda gets an update from the IT Manager that the system is now working, so she carries out a spot check with two employees who confirm that they are now able to submit their timesheets.

She updates the bulletin and settles back to daydreaming about the weekend.

But she knows come Monday, the incident and problem management review huddle will have a new talking point.

Major IT service incident

“Oh no!” Blake the NOC engineer exclaims.

Everyone turns towards him as he pointed at the main screen. Half the nodes are now flashing red—yikes. Sheryl, the NOC manager for this cloud provider, figures it’s either a core switch or hypervisor issue that’s affecting half of their clients’ virtual machines (VMs).

Blake logs the incident on their ITSM system, categorizing it as a major incident. Sheryl gets on the phone and sets up a conference with the cloud admins and the network administrators.

This will require all hands-on deck. The PR manager is looped in to the conference call, as he’ll need to inform clients and manage the coming social media storm.

The cloud admins soon realize that it was a bug on the hypervisor. They immediately reach out to the vendor by phone. To support this, the Cloud Admin lead raises a P1 ticket on their service portal.

By now, things are getting heated.

Calls flood the call center. The CEO is now involved, making personal calls to the leadership of the affected clients. The vendor wasn’t responding as quickly as possible, but the CTO is already two steps ahead and triggered the disaster recovery plan. The VM backups were spun on different servers and the incident was resolved in a few hours.

The following week, Sheryl would be seated at the problem management review meeting looking at feedback from the vendor as part of root cause activities. Her incident report featured heavily, and she foresees lots of changes ahead in order to ensure such a disruption does not again happen.

Incident management workflow & activities

You can see from these examples that any number of activities might help—or hurt—your attempt to address an incident.

In order to handle incidents in a way that meets the needs of customers and relevant stakeholders, your IT team will perform a variety of activities, generally in this order:

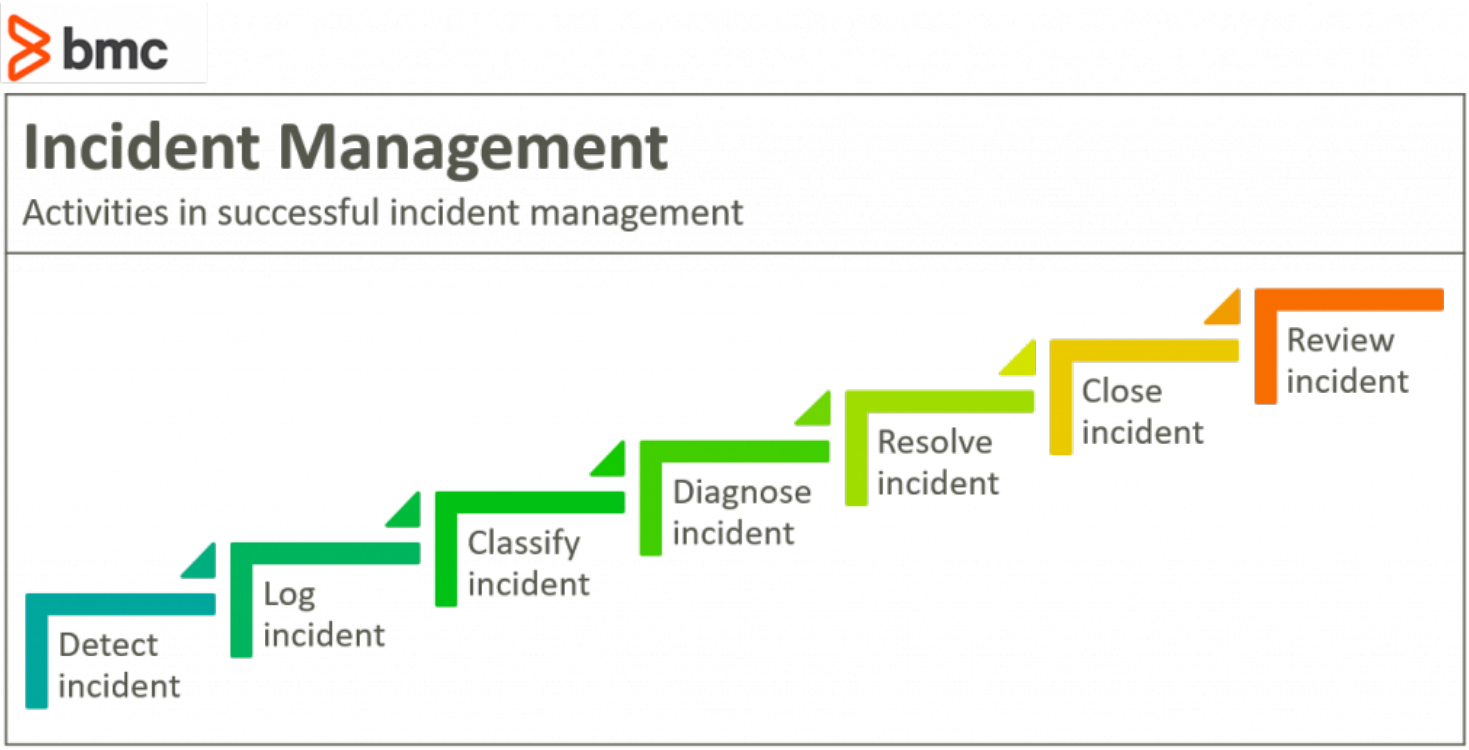

1. Detect the incident

Incident detection usually happens in one of two ways:

- A user reports a service issue and the service provider validates it as an incident.
- The service provider identifies an incident from alerts or trends from the components used to provide the service.

2. Log the incident

The service provider logs the incident. This should register it in a system for purposes of proper management, including:

- Assigning the right handler to the incident
- Tracking the handling progress, particularly the timelines

3. Classify the incident

In the incident classification phase, the service provider categorizes the incident in terms of:

- Type
- Impact, as in who and what is affected
- Urgency, or the speed required for resolution
- Priority, with regard to business and customer perspectives

Classification is useful for accelerating the process of identifying:

- Who should handle the incident
- What model, if any, is best suited
- Whether existing workarounds can be used

4. Diagnose the incident

During incident diagnosis, the service provider investigates in order:

- Identify what has gone wrong
- Determine the fastest way to recover normal service

Diagnosis can be done by one person (handler) where the symptoms relate to a previously known and documented incident. But, for more complex and/or relatively new incidents, a team of cross-functional representatives, known as a swarm, may conduct a joint investigation.

Diagnosis may result in an update to the classification of the incident.

5. Resolve the incident

Incident resolution refers to when the solution is applied—be it a workaround or a permanent fix. Resolution can take one or several forms:

- Implemented automatically
- Documented for the end user to apply it by themselves
- Handled by the support team
- Forwarded to a more skilled unit or even the vendor

Depending on the length of time the incident is taking and its classification, communication with affected users and stakeholders must be carried out in parallel, informing them of status and timelines.

If your resolution efforts are not bearing fruit at the required speed, you may need to backstep to diagnosis or trigger the disaster recovery plans.

6. Close the incident

Once the incident is resolved, formal incident closure of the record takes place. Closure might require:

- Communicating and confirming from users that the service experience is normalized
- Billing of handling activities
- Updating configuration information where required

7. Review the incident

During the incident review, sometimes known as an incident postmortem, the process owners or management may review how the incident was handled to determine what was done right and what went wrong. Both are useful in future incidents by illustrating what activities might need to be changed or reinforced.

Review can usher in process activities from other ITM practices such as:

- Problem management
- Service level management
- Information security management

- Release and deployment management
- Service design
- Change management (aka change enablement)
- Others as needed

Successful incident management: best practices

Speed is the name of the game when it comes to incident management. Customers, users, and stakeholders all want normal services to resume as quickly as possible, with the impact of the incident and its repeat probability minimized as much as possible.

For the most successful incident management, consider how your organization is set up for these factors:

- Detecting incidents early—and *before* customer impact
- Responding to and resolving incidents as quickly as possible
- Central managing of incident information in order to communicate, collaborate, and measure the incident response
- Ownership and coordination of incident handling activities
- Continual improving on all elements of incident management

There are many moving parts involved in incident management. Therefore, it is imperative that you apply a rigorous approach across all process activities, ensuring that service value and customer perception is not eroded by mishandling or poor coordination.

At the same time, continual review and analysis of incident management activities will ensure that a cost-effective approach, which maximizes on the service provider's capabilities, is maintained progressively.

Related reading

- BMC Service Management Blog
- How To Map the Incident Management Process
- 6 Ways To Improve Incident Management from The Service Desk
- The State of Incident Management Today
- Introduction to Critical Incident Response Time (CIRT): A Better Way to Measure Performance
- Digital Forensics and Incident Response (DFIR): An Introduction