

# Incident management: a getting started guide

With advice from real-life responders and planning templates for your team



# Table of contents

3	<b>Introduction</b>
5	<b>Chapter 01: Getting incident ready—defining key terms</b>
6	What's an incident anyway?
7	Capturing and identifying key fields
8	Workflow and status changes
9	Addressing service level agreements
11	Defining incident roles
13	<b>Chapter 02: Communication and collaboration</b>
14	How will internal teams communicate?
14	How will stakeholders and customers receive updates?
15	What should each communication include?
16	Top tips from real-life incident managers
18	<b>Chapter 03: Incident postmortem and practice run</b>
22	Resources
23	Appendix



# Introduction

IT incidents can cost between **\$2,300 to \$9,000 per minute** depending on your company's size and industry. That translates to over \$1 million for a two-hour outage. While lost revenue is a great reason to prevent incidents and improve response—money is not the only thing lost during a disruption. Other tangible costs include company reputation and employee morale, especially if devastating incidents are a frequent occurrence. It's arguably easier to recoup revenue than the good faith of your customers and employees.

With that said, incidents and downtime are as inevitable as the common cold. You can sneeze into your elbow, wash your hands, or live in a bubble—eventually you're going to get sick. However, if you take simple precautions like handwashing, and eating well—you can likely decrease the frequency and duration that you're ill. The same is true for incidents, eventually something will go down. But just as we have handwashing, medicines, and soups to protect our (human) health, we have strategies to protect the health of your services, systems, and infrastructure.

The key to reducing incident frequency and duration begins with the preparation. What constitutes an incident? Who will respond to that incident? How will customers be notified? Should some incidents be escalated sooner than others?

As your organization grows, an informal incident management process will not cut it. You can't protect millions in revenue and keep quality employees if you're still leveraging an ad-hoc phone tree in the year 2021. At the same time, setting up a formal incident management process can be intimidating. That's why we've created this guide.

We want to help teams like yours get started with a formal incident management process. In addition to easy-to-follow steps and templates we've included no-nonsense advice from real-life IT incident managers to inform your process. This guide is designed for those just beginning or scaling a formal incident management process.

#### **We'll cover:**

##### **Getting incident ready**

- ✓ Defining terms
- ✓ Capturing and identifying key fields
- ✓ Workflows/Status changes
- ✓ Service-level agreements
- ✓ Incident roles

##### **Communicating effectively**

- ✓ Stakeholder comms
- ✓ Internal comms
- ✓ Postmortems



# 01

---

Getting incident ready—  
defining key terms

## What's an incident anyway?

### Information technology infrastructure library (ITIL)

The most widely recognized framework for IT and digitally enabled services in the world. It provides comprehensive, practical, and proven guidance for establishing an effective [IT service management \(ITSM\)](#) system.

There's a bit of lead-up before creating an incident response plan. First you have to make sure your team is on the same page. To be successful, everyone needs to share the same high-level understanding of what incidents are, how they're currently managed, and what's not working. A good place to start is with the ITIL definitions of alert, incident, and major incident. Most likely if you're reading this guide you have familiarity with these terms. However, many organizations put their own spin on common terms, so it's a good measure for your team to discuss as a group and agree on what the terms mean for you.

### ITIL definitions of key incident management terms

An **alert** is a notification that a threshold has been reached, something has changed, or a failure has occurred. Some organizations use notification and alert interchangeably.

An incident is an unplanned interruption (or potential interruption) to or quality reduction of an IT service.

A **major incident** is the highest category of impact for an incident. A major incident results in significant disruption to the business.

[ITIL® glossary and abbreviations](#)

For example, imagine you're an online retailer. If the web app that your customers use to shop goes down and is unusable, that's likely a major incident that you need to respond to immediately. If your apps run on AWS, and there's a scheduled maintenance window coming up, that's an alert. The alert will require a response, but the impact of what an alert tells you is far smaller than an incident or major incident. To get everyone on the same page, you can use the table below as a model. We've included a blank one in the appendix for your use.

Term	ITIL definition	(Your company's) example
<b>Alert</b>	Notification that a threshold has been reached, something has changed, or a failure has occurred.	Scheduled maintenance during shopping hours for AWS EC2.
<b>Incident</b>	An unplanned interruption to or quality reduction of an IT service.	45 second web app outage, or 100 users can't access their account information.
<b>Major Incident</b>	The highest category of impact for an incident. A major incident results in significant disruption to the business.	Web app outage for more than five minutes.

## Capturing and identifying key fields

Once you've agreed on the definitions of key terms, the next thing you'll want to do is determine which fields should be required for incidents. Determining which fields you want to capture is important for the reporting after the incident. No matter which ITSM tooling you're using there will likely be many fields that you can include. Some key examples are: priority, impact, urgency, reported by, assigned to. Secondary examples include: time to first response, time to resolution, time incident began, time incident closed, components, services, related services, and more.

## Workflow/Status changes

Within whatever tooling you're using, there will be different workflows and statuses that trigger an incident. Typically when something goes wrong, an alert will get created, and then trigger an incident based on the rules that you set. Once the incident is created, responders will be notified and begin working to resolve the incident. There's a careful balance between customizing your toolkit and over-complicating your use case. Keeping the below tips in mind will help you strike a balance.

When setting up the logic for incidents some things to consider and discuss include:

- 1 Whether the alert priority and the incident priority should always match.**  
For example the alert might be a P1, but the incident itself may be a P3. This is highly dependent on your individual use case, but it's worth discussing.
- 2 What parameters change a notification from an "alert" to an "incident."**  
Factors can include the number of alerts coming from a single source, the amount of time that the alert is true, and more. For example, if Amazon CloudWatch is operating at 90% CPU for 2 minutes, this is likely an alert, however if it is operating at 90% CPU for a half hour you might want it to trigger an incident workflow. Leverage what's already built into the tool before planning your own special circumstances.
- 3 Failsafes for presenting an unnecessary incident from triggering.**  
Some companies prefer human intervention before an incident is declared, others prefer to lean on automation to reduce mean time to resolve. What works best for your organization is dependent on your use case, but be sure to discuss as a group so that everyone is on the same page.

**“** When we changed our ITSM system a few years ago, instead of changing the way everyone works and leveraging what the tool had in place, we customized the new system as much as we did the old system. There is no end to the headaches this causes. If you're just starting I would stay away from overly-complicated customizations—use what's there to the best of your ability.

Michael Marques, ITIL Certified ITSM Incident Manager, Bose



The point to stress here is don't reinvent the wheel. Choose the best options for your use case, but also try to leverage the existing workflows of the tooling you're using. Overcomplicated customizations will weigh down the team and make communicating more confusing. Once you've socialized incident and alert parameters, you can move on to the next step of defining service-level agreements (SLAs).

## Addressing service level agreements

### Service-level agreement (SLA)

An agreement between an IT service provider and a customer. It describes the IT service, documents service level targets, and specifies the responsibilities of the IT service provider and the customer.

[ITIL® glossary and abbreviations](#)

SLAs can apply to responses, incident and alert resolutions, and more. Different services, products, or customers might have different SLAs. For example, the CEO of a major client might have a guarantee that any issue she reports is responded to within 30 minutes. You might set an internal SLA for a business-critical service, like your company's customer-facing web app.

As you're sitting down and defining SLAs, communicate with the different stakeholder and customer groups to get a sense of what they typically expect. If you have an online community, or a beta testing group, or even partners or high-profile clients that are invested in your success, it would be great to create a program and run the proposed SLA's by them for feedback. Also be sure to evaluate what other competitors in your space offer to check and balance customer expectations. If you're hungry and ordering a pizza, and Store A will deliver in 20 minutes, but Store B delivers in 45, you'll likely give Store A your business. Folks in need of support are desperate, just like hungry people in need of pizza—keep that in mind for all of your customer-facing communications.

“ We took a baseline by looking at tickets from two years previous and those tickets’ priorities. Then, we looked at the average time to resolve those tickets and determined a reasonable turnaround time. After settling the SLAs, years later we had a workshop with our customers to see how our SLAs were aligning with their expectations. We leveraged some Lean tools to get the voice of the customer. This is something I highly recommend.

Michael Marques, ITIL Certified ITSM Incident Manager, Bose

Once you set SLA’s be sure to socialize them both internally and externally (where applicable). Most ITSM solutions enable you to track and label SLAs within issues or incidents. This helps to manage expectations for both the team working the ticket, and the customer waiting for the solution. Knowing when they’ll hear back on their problem can reduce the anxiety and frustration of the customer while they wait. This can be done within an auto response to the ticket, notices on the submission form, and a list of SLAs on your request portal.

	Low	Medium	High	Critical
<b>Priority description</b>	Little to no effect on the ability to do one’s job.	Limited loss of normal functionality.	Loss of normal functionality.	Severe disruption or degradation.
<b>Example</b>	Customer is a graphic designer and the request is for access to Spotify.	Customer can access email via web browser, but not directly via the email application.	Customer can’t access their account profile.	Retail store website is down.
<b>Urgency</b>	Low	Medium	High	Critical
<b>SLA target</b>	48-72 hours	8-12 hours	4-6 hours	2 hours

Once you’ve determined how quickly different types of events should be resolved, the next step is to define incident roles.

“ Everything starts and end with processes and expectations you set for communications. One of the best changes we made was adding an SLA for the time we expect a technician to communicate back to the customer.

Michael Marques, ITIL Certified ITSM Incident Manager, Bose

# Defining incident roles

Setting incident roles before an incident ensures organization during the chaotic moments when everything breaks. The roles you need to fill can vary by incident, organization, or team. For the purposes of this whitepaper we'll keep it simple. Leveraging the right roles for your situation helps incident response efforts to run smoothly. For example, it's key to have someone who is not actively working on the incident to handle the communication to management and stakeholders.

This enables everyone to focus on their responsibilities and prevents interruption to the flow of information. During a massive fire, firefighters aren't talking to the press—they usually have a media liaison for that. This way, those most qualified to put the fire out can focus on the problem at hand, and the media liaison can inform the public. IT incidents are no different, leave the responders to the fix and have a communication officer handle stakeholder outreach.

Here's a basic list of incident response roles to work off of:

<b>Incident commander</b>	Responsible for managing the incident response process and providing direction to the responder teams.
<b>Communications officer</b>	Responsible for handling communications with the stakeholders and responders.
<b>Scribe/Note taker</b>	Responsible for documenting information related to the incident and its response process.
<b>Subject matter expert</b>	Technical domain experts who support the incident commander in incident resolution.

“ Let the responders do their incident-related jobs. Responding to frantic managers and customers AND trying to fix something takes more time than just working on the problem. Usually poor user and stakeholder experiences are related to poor communication, this is why someone on the response team needs to be dedicated to communication.

Patricia Francezi, Jira Admin Service Manager, iDev

If you're unfamiliar with incident management, scribe might seem like an odd role to include. However, it's one of the most valuable. The scribe could be your service desk agent or whoever is responsible for keeping the incident record updated. One of the best tools you can have in your toolbox are detailed, clear notes. You need to know what was changed, the order it was changed in, the impact of each change, and which teams completed each change. Just like it's hard to be the communicator if you're the one working on the incident it's also hard to take detailed notes—which is why it's important to have a detailed record.

For more roles and use cases, this [page](#) is a great reference and explains various approaches to incident management roles. Now that we've covered roles and their functions let's go into some tips on communicating effectively.



Take really good notes, try to find out what was happening when the incident happens, from all involved parties/ departments even things that seem like they wouldn't be related. Just put in the notes with a time stamp. It's interesting what things will just fall out when you get it all on the timeline.

Kimberly Deal  
Information Security Manager  
Senior Jira Administrator  
Wells Fargo



Change only one thing at a time. Write down everything that you change and the results of each change.

Matt Doar  
Senior Jira Administrator  
LinkedIn



# 02

---

## Communication and collaboration

# Communication and collaboration

Communication during an incident is a big needle-mover. Poor communication can lead to frustration, longer time to resolve, and unhappy customers. When setting the bar for communication take into account the following:

## How will internal teams communicate?

Within your organization talk about what will be used as the main method of communication during an incident. Whether you choose a ChatOps tool like Slack or Microsoft Teams, a phone bridge, or a video conference, ensuring that everyone knows where to communicate keeps collaboration streamlined and organized and reduces chaos on the day of.

## How will stakeholders and customers receive updates?

In most cases, organizations let external stakeholders know of an incident by using a public-facing status page. In the case of a severe outage, proactively emailing customers might make sense as well. No matter what methods you choose be consistent and set the expectations. Let your customers and internal teams know where they should look for updates. Also think ahead, incidents increase support ticket and call volume. Simple touches like adding a web banner during an incident to your request portal or status page can clue the customer in and reduce the burden placed on support.



The best advice I could give is communication is the key, there is nothing more important. Everything starts and ends with processes and expectations you set for communications. First and foremost you should set the standards for communication, when you'll communicate and how. Keep the human element of helping people.

Michael Marques  
ITIL Certified ITSM Incident Manager  
Bose

Whether internal or external, communication with stakeholders and customers needs to be consistent, clear, and frequent. Which brings us to the next topic, what “good communication” looks like.

## What should each communication include?

We’ve all gotten a useless update when one of our favorite tools is down. An update like “Service is down. We are currently investigating.” is not sufficient. The communication should explain and describe what specifically is happening, when it was detected, who it is affecting, what can be expected as a result, and when the next update will occur. This goes back to the managing expectations piece, in a stressful situation out of a person’s control—knowing what to expect can take a heap of stress off their shoulders.

### Here’s an example of a helpful communication:

*November 11, 2020 05:45 a.m. UTC*

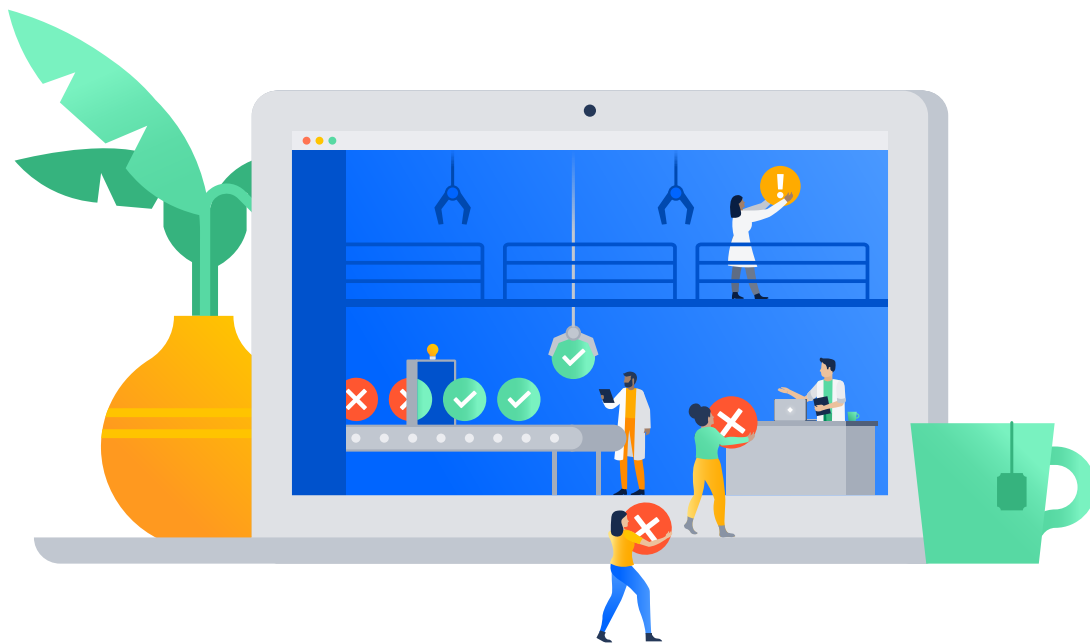
Service is currently down for North American customers. The issue was first detected at 05:30 this morning. Our team is aware and looking into the problem as well as working toward restoration. We will post the next update at 06:00. As a result of the outage customers are unable to access their profiles, don’t hesitate to reach out to support with questions.

Let’s go over why this is helpful: you know what is down, you know what to expect as a result, when it started, and when the next update is coming. This information is vital to clients and customers who rely on the products and services you provide. When writing communications don’t promise immediate restoration and don’t tell your customers that the problem is resolved before it’s fully confirmed. In the beginning moments of an incident sometimes the full impact is still unclear.

## Top tips from real-life incident managers

“ Have separate people to handle communications. And all communications should say when the next communication will happen.

Matt Doar  
Senior Jira Administrator  
LinkedIn



“ Don't promise time to restore if it's a big or a complex issue because if you fail to fix it in that time the pressure will double. Instead, promise communication about status and progress so the techs can work the issue, and the stakeholders have peace (or at least be less angry).

Patricia Francezi  
Jira Admin Service Manager  
iDev




The table below is a great discussion piece to use when going over what the communication standards and expectations are with your team and within the organization as a whole. Although every company operates differently, it's prudent to involve internal and external communications teams (including social media, public relations etc), engineering, IT operations, dev teams, support teams, leadership teams, and a focus group of investors (if applicable). Involving all these groups may seem a bit much, but these are the groups that will be put in the hot seat if communication goes poorly. If something is broken, a customer calling into support does not care which team is responsible—the support agent is stuck dealing with their frustration. When an angry customer tweets a complaint at a company during an outage, the social media or public relations team takes care of it. Enabling everyone to participate in the discussion ensures that the company is aligned and empowered to communicate and reduce the stress that surrounds a chaotic incident. It also enables these teams to plan ahead and create templates for such events.

Type of communication	Standard
<b>Customer/External communications</b>	<ul style="list-style-type: none"> <li>• Determine the main channel of communication (i.e. a status page, Twitter, a website banner, a notice on the request portal).</li> <li>• Determine the time interval of each update.</li> <li>• Ensure that every update indicates when the next update will occur.</li> <li>• Be explicit.</li> </ul>
<b>Internal communications</b>	<ul style="list-style-type: none"> <li>• Determine the main channel of communication. (Slack, Microsoft Teams, phone bridge, video conference, etc)</li> <li>• Be specific about what exactly is wrong, and the specific steps being taken to resolve.</li> <li>• Determine an interval for regular updates.</li> <li>• When the incident or problem is resolved, clearly communicate how it was resolved, and how it was verified as resolved.</li> </ul>

Being detailed doesn't just apply to stakeholder communications and status page updates however. When resolving a problem or fixing a bug be sure to explain exactly what was done and how it has been verified as resolved. A resolved or closed incident with no details doesn't give confidence to other team members about an incident's status.

Once you've discussed incidents, incident roles, SLAs, and communication standards, it's time to put it to the test. Get everyone together for a practice run!



# 03

---

## Incident & postmortem practice run


# Incident & postmortem practice run

When you change or introduce a new process the best way to detect flaws and socialize that process is to run through it “live.” You don’t want to find out in the middle of a real incident that the method of communication you’ve chosen doesn’t work or that the Incident Commander doesn’t understand the scope of their role.

After the process has been shared and socialized, send a calendar invite to your team to block off time to fake the incident. Test your alerting, on-call schedule notifications, etc. Behave as if it is truly an incident and walk through a fake resolution. Assign other folks outside of your team to play external stakeholders (this helps remove bias from the communication quality). The goal is to check the following:

- 1 Did everyone know what to do?
- 2 Did everyone understand their roles?
- 3 Was the internal communication clear?
- 4 Did external stakeholders feel informed?
- 5 Did the process work well for the team?

Leverage the postmortem process to uncover any glitches, record lessons learned, and gain insight from everyone on the team. Keep in mind that postmortems aren’t about blame, they are meant to celebrate successes and improve future incidents. The postmortem process should be positive and treat any identified problems as an opportunity.



“

Plan for the disasters.  
Test the plans. Assume  
the worst will happen.  
Don’t Panic!

Matt Doar  
Senior Jira Administrator  
[LinkedIn](#)



Don't name names, throw blame, or call out people on the incident call. Present the problem, present the symptoms, offer solutions, and be a positive influence. This is a rough patch for everyone involved.

Kimberly Deal  
Information Security Manager  
Senior Jira Administrator  
Wells Fargo

Once the postmortem is complete and everyone debriefed there is still one more thing to do. Share your findings with leadership and key decision-makers. Inform them of the new process and walk through the postmortem learnings for the fake incident. Without leadership buy-in and support it will be difficult to get folks onboard. If leadership helps champion the new process, the team will be more motivated to make it work.



One more thing for major incidents. The most important step is getting leadership buy-in and support. Without the support of leadership holding people accountable to root cause analysis, reporting and debriefs are impossible. If leadership doesn't make it a priority the people who are expected to do the work will not make it a priority.

Michael Marques  
ITIL Certified ITSM Incident Manager  
Bose

Now that we've taken you through all of the high-level steps of incident management, you should be officially ready to start planning an incident management process. Be sure to check out the resources section to learn more on each aspect of incident management and response. You can also refer to the appendix for helpful planning and discussion tools.



---

## Resources and appendix

# Resources

- ✓ [The Atlassian Incident Management Handbook](#)
- ✓ [Incident communication template generator](#)
- ✓ [All about postmortems](#)
- ✓ [All about incident management](#)



# Appendix

## Defining key terms worksheet


Term	ITIL definition	(Your company's) example
<b>Alert</b>	Notification that a threshold has been reached, something has changed, or a failure has occurred.	
<b>Incident</b>	An unplanned interruption to or quality reduction of an IT service.	
<b>Major incident</b>	The highest category of impact for an incident. A major incident results in significant disruption to the business.	

## SLA table

	Low	Medium	High
<b>Priority description</b>			
<b>Urgency</b>			
<b>SLA target</b>			

## Communication worksheet

Type of communication	Standard
Customer/External communications	<ul style="list-style-type: none"><li>• Determine the main channel of communication (i.e. a status page, Twitter, a website banner, a notice on the request portal).</li><li>• Determine the time interval of each update.</li><li>• Ensure that every update indicates when the next update will occur.</li><li>• Be explicit.</li></ul>
Internal communications	<ul style="list-style-type: none"><li>• Determine the main channel of communication (Slack, Microsoft Teams, phone bridge, video conference, etc).</li><li>• Be specific about what exactly is wrong, and the specific steps being taken to resolve.</li><li>• Determine an interval for regular updates.</li><li>• When the incident or problem is resolved, clearly communicate how it was resolved, and how it was verified as resolved.</li></ul>





# Getting incident ready

## Basic Checklist

- ☐ Define and socialize alerts, incidents, and major incidents.
- ☐ Set and share SLAs
- ☐ Define Incident roles
- ☐ Set channel and standards for stakeholder comms
- ☐ Set expectations for internal comms
- ☐ Simulate an incident
- ☐ Simulate a postmortem
- ☐ Socialize incident management and response plan with upper-management
- ☐ Never stop improving, respond, resolve and learn from every incident

