Incident Response: 6 Best Practices For The Modern Enterprise

Modern organizations today are managing increasingly complex technology portfolios and pressured to deliver on innovation—all while facing far higher stakes than ever before when it comes to maintaining service performance and reliability. While these demands may seem like a paradox, many organizations have been successful in implementing processes that enable them to balance both agility and risk. In this post, I'll touch on the importance of integrating incident response with your ITSM tool and walk you through the steps on how to effectively balance agility and risk.

Step 1: Integrate Incident Response and ITSM

You can't add minutes during an outage, so prioritizing your planned work outside of an incident effectively is key—and part of that is using an incident resolution platform like PagerDuty to manage and tie your unplanned work back to the planned work that's tracked in your ITSM tool like Jira, ServiceNow, or Remedy.

How does that help? First, information flows from ITSM into PagerDuty so that responders know what has changed and who is reporting an impact. Next, follow-up items from PagerDuty are sent back into ITSM, including outcomes of the postmortem that need to be prioritized.

A given employee may have dozens of prioritized tickets in an ITSM tool, but they should only ever have 1 (or ideally 0) assigned to them in PagerDuty at a given time so they can focus on customer-impacting issues that require immediate responses. Similarly the concept of unassigned incidents doesn't exist in PagerDuty—if there's a problem, someone is responsible for that problem.

Step 2: Proactive Mobilization

Simply put, the easiest way to speed up your response is to start it earlier. The best way to do this is not tracking what affects your machines, but what affects your customers. Organizations that use Real User Monitoring can track whether users are able to successfully load, download, or buy their tools. Additionally, since you're primarily looking to detect problems before they affect users (although at the cost of some false positives), monitoring the underlying infrastructure is equally important to identify the cause of a customer-facing problem.

Automation also plays a role in speeding up incident response, and your monitoring tool should automatically assign problems to an owner. Along those same lines, to prevent an issue from affecting your revenue, the monitoring tool should also assign and immediately notify someone about all issues above a certain priority using that person's preferred communication method (phone, email, SMS, etc.).

To make automation easier, PagerDuty integrates with hundreds of monitoring tools. So, for example, if your monitoring tool detects that your shopping cart has gone from slow to completely non-responsive, PagerDuty can automatically create an incident with the correct priority to ensure the responder has all the information.

In the same vein, create automated workflows whenever possible. If a Sev1 needs to pull in executive stakeholders, automate that response play.

Step 3: Have a Defined Process

Remove ambiguity, confusion, and wasted time during a response by defining your process and clarifying the different roles involved. We recommend including the following roles: Incident Commander + Deputy, Scribe, Customer Liaison, and Subject Matter Experts. (For more details as to what each role means, visit Different Roles.)

During an outage, things can become a bit of a madhouse and the organizational hierarchy takes a backseat to the response roles. When executives start to randomize the defined process, you need to remove them from the process and communicate clearly and concisely why certain processes are followed—and if the CEO wants to change the process on-the-fly, they can decide to become the Incident Commander.

To help everyone keep it together, remember the following:

- **Poll for strong objections.** Ask for objections, not consensus. Doing so ensures you don't get stuck waiting for non-urgent discussions and consensus-building instead of acting to resolve the issue.
- **Time-box and assign tasks to individuals.** A lot of information will be coming in during an incident and clear, concise communication is key during times of crisis. Assigning tasks and time limits to tasks helps each role focus on one thing, reducing confusion and double work—and ideally, time to resolution.
- **Standardize lingo and etiquette.** Ensure everyone knows when and how they can speak up. Keeping the tone and discussion practical and focused on the issue, without emotion, is key to effective communication and response.

Step 4: Build Your Communication Strategy

It's important to define a process around communication to people outside of the core response team as well. Depending on the type of incident, you could be dealing with internal customers (we often call them stakeholders), external customers, and even the market at large. For instance, when responding to a security incident, you may need to loop in the legal department in addition to other executives.

These groups all need to be kept up to speed on an as-needed basis, but the wrong place to do that is where the responders are working. The last thing you want is someone joining the call and asking for a status update as this disrupts the people trying to discuss fixes during the call. To my point earlier, you don't want an executive getting on a call and demanding that the team fix the outage in 10 minutes. This implies the team is not already working as quickly as they can. It's demotivating and doesn't contribute anything helpful for the response. This is where the Customer Liaison comes in—using a feature like PagerDuty's Stakeholder Engagement, the Customer Liaison can provide streamlined, real-time updates to relevant stakeholders across the business.

Here are a few other ways to improve real-time communications:

- **Have a conference bridge for internal discussion.** Humans are social animals and this seems to be the most natural format. Use the conference call tool that your users are already familiar with—an outage is not the time to learn a new tool. Automatically attach the conference call information for major incidents.
- **Have a chat room for logging actions.** This gives those jumping into a response an ability to get up to speed without asking repetitive questions and provides a time-stamped record of the response. Additionally, many companies are starting to trigger response actions directly from bots in the chat room.
- **Provide proactive, scheduled updates for your stakeholders.** Set up an incident status page so they can stay up-to-date on relevant, real-time information. This is essential to prevent stakeholders' urges to step in and become roadblocks.
- **Determine notifications ahead of time.** Decide what criteria and timeline responders should use to notify your stakeholders, customers, or downstream users.

Step 5: Postmortems

Postmortems are how you fix a long-term problem. They give closure to people after a particularly stressful event and guarantee that your team can take well thought–out and productive action on some of the immediate patches you made in the heat of the moment to solve a problem.

So what does an effective postmortem look like? It should:

- **Focus on prevention and learning.** Your team is looking to understand what can be changed to avoid this issue in the future.
- **Be transparent, blameless, and apolitical.** The goal is to get all the relevant information, and the last thing you want to do is foment grudges. Blame impedes information flow. The only acceptable blame is if you've uncovered an intentionally malicious employee, which is exceedingly rare.
- **Be oriented around improvement.** This applies to both the system's resiliency and the response process. The goal is to always get better.
- **Target a root cause.** We find the "five why's" helpful here.

- **Be required for major incidents and streamlined to save time.** No one wants to do postmortems, but they are an essential tool to maximize the impact of your planned work. To make them easier, we've built an integrated postmortem tool modeled on our customers' existing processes. It can save you hours toggling between tools to collate information, as it automatically creates a timeline with relevant PagerDuty and chat activity.

We post all of our postmortems internally using our postmortem tool. We view postmortems not only as learning for our team, but also as an input to our best practices training, where we share our experiences and learnings with our customers.

For more postmortem tips, download our detailed e-book.

Step 6: Training and Practice

You can't expect your incident response process to be fantastic if you only use it every once in a while. Not every service fails often and some people get more practice than others. But everyone should be practiced so that when something does happen, you and your team are ready.

- **Make shadowing and onboarding easy.** A solution like PagerDuty makes it easy for overwhelmed responders to pull in help. One of our braver customers starts everyone on call solo—if a new hire can't figure it out from the runbook, they can add their mentor as a responder, and over time, the percentage of incidents they need help with drops.
- **Record your outages to use for training.** These recordings are a gold mine and help teams understand what actually happens in real failure scenarios. They are also useful for postmortems.
- **Pre-mortems ("If this breaks, what would I look for?") are valuable as a training exercise.** They can also help identify places where you can add additional monitoring for root causes or pre-emptive warnings. For instance, if checking the database connection is the first thing you would do if you were alerted because your e-commerce was down, then set up monitoring on that connection and send that data into PagerDuty—even if the app isn't affected.
- **Implement Failure Fridays.** Chaos engineering is probably beyond most organizations at this time, but we get a lot of mileage out of our Failure Fridays, such as uncovering implementation issues that reduce our resiliency and proactively discovering deficiencies to prevent them from becoming the root cause of future outages.

The less time you need to spend fixing unplanned outages, the better your services are, which results in happier customers since customer-impacting incidents are likely the worst thing that can happen to a business. They damage brand reputation, cause huge losses in customers and revenue, inhibit employee productivity, and slow down morale, among other things. If you can get to a point where you are as efficient as possible and are able to respond to major incidents without chaos and stress—with the attitude that you will learn and improve from each one—you will achieve a winning and empowering culture that stands to delight both your customers and employees.

Interested in learning more about incident response? Check out our incident response documentation page.