PagerDuty

Choosing an Incident Management Platform

What to look for and what to avoid

The search for the ideal solution

How critical is it to automate incident management processes at your company?

Today, just minutes of downtime can mean thousands of dollars lost, negative customer sentiment, or hundreds of lost customers. These consequences mean IT organizations must maintain and manage cloud and on-premises infrastructure, applications, APIs, and containers—all while rolling out enhancements and upgrades in near real time to meet ever-changing customer demands. Additionally, IT organizations need to figure out a way of blending traditional IT service management (ITSM) models with more Agile, DevOps, and SRE best practices.

There are many tools and solutions to choose from when it comes to accelerating your incident management processes, and they all have varying features. To remain competitive, it's critical that you use an incident management platform that provides real-time visibility into an issue, along with full context, when something goes wrong.

The ideal platform should:

- **Decrease and prevent downtime**
- **Adapt to non-linear and ad-hoc processes that arise with real-time work**
- **Give teams situational awareness around their operational health, responder health, and infrastructure**
- **Help deliver a reliable product to your customers**

This ebook summarizes the top 8 questions organizations should keep in mind when choosing a modern incident management platform:

1. Is the incident management platform always available?
2. How are stakeholders updated and how often?
3. Can you have a service-based incident management setup to provide business context and organizational flexibility?
4. Do you have the flexibility to work where you want?
5. Do you have real-time visibility into incidents and the health of your applications?
6. Does the platform filter out noise so responders can focus on actionable alerts?
7. How engaged is the user community?
8. Does it provide situational awareness and real-time data that responders can act on?

1 How can you ensure your incident management platform is up when you are down?

Does your incident management platform use maintenance windows?

Customers expect around-the-clock access to bank account information, to buy groceries, and even to request a rideshare service. This means the onus is on companies to ensure their digital businesses are up and available at all times—or they risk losing customers and revenue.

With these heightened expectations, maintenance windows—where teams had to protect expensive hardware and prevent software from failing—are a relic of the past. The rise of cloud architectures, containers, and SaaS has allowed for more flexibility and innovation, and most modern companies have embraced DevOps culture and agile development and have either stopped having maintenance windows or are phasing them out.

A good incident management platform should reflect this new reality and foster an agile way of working. If not, what happens if you have an outage while your solution provider has their maintenance window? Oftentimes, it means missed alerts, extended downtime, lost revenue, and unhappy customers.

BEST PRACTICE

Look for an incident management platform provider that has phased out maintenance windows to ensure you don't miss any alerts.

The importance of SLA transparency

Speaking of alerts, the solution you're trusting to alert you about your outages also needs to be transparent about its own outages. Making its security and reliability measures public should be a top priority for any incident management provider because doing so provides customers with more insight into how much the vendor values security and reliability.

The first step is to look for clearly worded service-level agreements (SLA). An SLA is an agreement between a service provider and client about the commitment and expectations of the provided service, and it should be easy to understand. Providers that make SLAs complicated or use vague language can cause confusion between parties as to what the expected terms actually are.

Additionally, you should ask the provider you're considering whether they keep a public record of their postmortems. Questions you can ask

include: Do they have a public-facing status page? Do they post about their outages on social media so both they and industry practitioners can learn from them together? Do they share their response process?

Asking the above questions helps ensure that you, a potential customer, will always have full visibility into your platform provider's status, as well as a clear understanding of how they're resolving issues and continuously improving resolution processes on the back end.

BEST PRACTICE

Make sure your incident management provider has a very clear, all-encompassing SLA to ensure that the service promised is the service delivered.

Do they enforce API limits or throttling?

In order to respond effectively to incidents, you need to ensure that your incident management platform can receive and process all alerts. Since every business today is a digital business, and the overall complexity of the technical services that support businesses continue to increase, more and more alerts are being created—which also means that alert storms are a common occurrence. With all these alerts coming through, companies need more information than just how many alerts or API calls they send or receive per day. They also need to know which alerts they should be focusing on and have access to full context in order to properly assess the severity and business impact of any issue.

Your platform of choice should be able to increase capacity to handle an increase of up to 10× API calls per day. A platform that doesn't have this ability to scale means you could miss important alerts, learnings from the data coming from your infrastructure, and the context you need to effectively troubleshoot and diagnose an issue. In short, if your platform limits the number of API calls per day, you risk missing alerts that could seriously impact both your business and your customers.

BEST PRACTICE

Ensure your incident management platform doesn't use complex calculations to throttle API input—if it is, it's protecting itself, not the customer (i.e., you).

2 The importance of stakeholder communications

When a major business- and customer-impacting incident occurs, technical responders aren't the only ones who need to take action. Other stakeholders from across the company—technical and non-technical—need to mobilize as well.

These "secondary responders" need to be kept up-to-date with incident resolution progress so that they can, for example, take

action to help mitigate negative business impact, such as putting together media talking points to do damage control. For employees in customer-facing roles, they need to understand how an incident will impact the business to inform their decisions as to how and what to communicate to customers. The ideal incident management response platform should enable response teams to automatically share concise and actionable status updates with those who need to know, and use various communication methods to do so.

For instance, sending notifications on specific incidents to stakeholders is one way to keep them informed. But what if said stakeholders hear it from a customer first and haven't received any information from headquarters? In this case, having access to a status dashboard that displays the health of pre-selected business services can help employees understand the current health of systems at a glance, review what has happened historically, and view any upcoming service changes like maintenance and system upgrades.

BEST PRACTICE

Look for an incident management platform that provides various methods of communication with stakeholders across the organization.

3 Does your incident management platform allow you to shift from a team-based approach to a service-based one?

You likely set up an incident management process at your organization because you need a team or someone to respond quickly when an incident arises.
But have you optimized your setup?

At PagerDuty, we've seen many companies use other solutions to set up their configuration with a team-first approach because it's fast and easy to create an on-call rotation and ensure everyone is on that rotation. But if you take a step back, a more optimized approach is to think about the services your teams support; after all, services are built to last, and they typically outlive the teams that originally developed them—people come and go, teams organize and reorganize all the time, and people take on new services and inherit old ones.

Taking a service-based approach to incident management setup allows you more flexibility with regard to changes. For instance, if a team's make-up changes over time (e.g., new people join, other people leave), the services they support can remain the same, which means less time is spent rearchitecting your incident management process whenever there's a reorg.

What exactly is a service-based approach?

A service-based approach requires you to identify your top-level business services that are distinct parts of the products or applications your customer interacts with to perform tasks. For example, "login," "shopping cart," and "search" are all considered business services.

Then, for each business service, identify technical services that contribute to that business service. Each technical service should ideally be owned and developed by one team at a time, even if multiple teams contribute to maintaining it long term.

Once you've identified your business services and the corresponding technical services that support them, you can now do a lot of interesting things. For example, teams can now see what's happening in real time across the business to better understand if an issue is isolated or has a broader impact, allowing for better coordinated response when it does span multiple teams and services.

When set up correctly, each disparate business service ties directly to a customer-impacting event, allowing you to calculate the true cost of an outage and connect that cost to the underlying technical service.

The benefits of a service-based approach

There are many reasons why you should orient your incident management and configuration setup around services first. Potential benefits include:

- Increased visibility to better understand the health of services and improve both internal processes and root cause analysis
- Detailed insights to see how services are trending, even for services that are still in the "normal" operational range, in order to identify hotspots and predict outages before they occur
- The ability to easily and quickly see which team supports what services vs. first having to go through multiple teams and understanding those layers before getting to the service

BEST PRACTICE

Use a service-based approach to set up your teams and services from the very beginning to better understand cost, as well as failure risk on a service-by-service basis.

4 Can your teams work in the tools they prefer?

In a given day, we use many applications to get work done. When it comes to incident response, we need to focus on the best ways to communicate and manage ad-hoc workflows, and it's imperative that teams work in the application that best enable them to decrease response times.

Can your teams work in the interface they're comfortable with, such as ServiceNow, Slack, Microsoft Teams, Jira, or on mobile? Change is hard, so being able to work the way they want ensures easy adoption, adherence to current processes, and allows teams to communicate in a manner that they're used to. It also ensures the right data is captured for postmortems. Another thing to consider is that incidents don't always happen while teams are in the office. In order to reduce resolution times, an effective incident management platform should be able to integrate with a number of different tools so responders can run an entire incident process—no matter where they are. Therefore, the integrations offered by the platform are important because the more integrations it has, the more likely users can work in the tools they prefer, which offers:

- **Convenience.** Users don't need to learn a new system or UI, which reduces training time and cost.
- **Speed.** Users don't need to switch between systems to respond to different incidents, so they're able to respond faster to all incidents.
- **Improved data accuracy.** Humans make mistakes—and forcing users to manually update multiple tools leads to inconsistent and missing data.
- **Automation.** Ensure triage data is synced to a system of record like SNOW or JIRA so no details are lost.

Does your team have the flexibility to run a full incident response process from anywhere, anytime?

Integrations are just one aspect of being able to "work where you are." The ideal incident management platform should offer the ability to run an incident management process from a mobile device. This means more than just acknowledging an incident; users should be able to create tickets, run response plays, notify stakeholders, run remedial actions such as rebooting servers, and get all the context they need to fix issues—all from their mobile device.

BEST PRACTICE

To improve incident management processes and shorten resolution times, find a platform that enables teams to work where they want and in the tools they prefer.

Does the platform provider have a well-documented REST API?

A platform with a well-documented public REST API will allow you to integrate with whatever monitoring tools you need to continuously deliver a high-performing, dependable customer experience, and it will fit into your environment however you need it—both today and in the future. A REST API will also allow your engineers to customize alerts, incident behavior, and workflows to their liking.

Additionally, without a public API, there's no way to insert or extract object data and easily customize your environment to fit your exact needs around data and workflows. Some good questions to ask that can determine whether this is important to you are: What monitoring, ticketing, deployment, and collaboration tools are you using today? Do you or will you ever have any custom scripts monitoring your environment?

IT toolchains are regularly evolving and expanding due to more innovative solutions coming to market than ever before, and you should ensure the tools you invest in can complement the changes. If the ultimate goal is improved efficiency, then that means creating a hub for all of your events across APM, logs, health checks, tracing, error alerts, tickets, deploys, and more.

BEST PRACTICE

Using a platform with a well-documented REST API will provide your teams with the flexibility to easily customize alert incident behavior and workflows to their liking.

5 How do you reduce the number of alerts—automatically, not manually?

Operational noise is growing at an exponential rate, with millions of events taking place across an organization's systems every day. But what's your capacity to deal with that noise? How do you know if an alert is actionable? Writing complex rules and exception handling will only take you so far and will ultimately be detrimental in the long run.

As the number of incidents increases, you're likely not throwing more people at them; instead, your teams are likely trying to get more efficient at solving (and preventing) them. A good incident management platform should make use of machine learning that meets the needs of modern, agile teams. It should be built on a few key tenets: ease of use, democratized access, and machine learning that provides clear insights and continuously improves as more data is fed into the system.

Machine learning is a new approach to incident management that makes it easier for organizations to aggregate their monitoring data, automate common workflows, suppress noise, and give responders complete situational awareness. Ideally, the machine algorithm should observe how humans interact with the platform and use that information to change how incidents are handled. Many options are available, such as grouping incidents together, suppressing incidents, changing routing, or suggesting custom actions like restarting a server.

The advantage of using machine learning over rules is that it requires less set-up time and manual maintenance since you don't have to create rules that constantly need updating. It can also be more accurate as it takes into account real-world situations and human actions that may not have been considered during setup and planning.

BEST PRACTICE

Find a platform that utilizes machine learning to observe human behavior and combine that with system data to optimize how incidents are handled and provide suggestions on how to manage an incident.

6 Does your platform provide real-time situational awareness into incidents and overall health of your critical applications?

During a major incident, the last thing your responders have time for is digging around for information from past incidents. Your incident management platform should provide a comprehensive view of application and service infrastructure in real time, context about what's currently happening, and how similar incidents were resolved in the past.

After all, most major incidents aren't "cord-cutting" events—in other words, systems don't just suddenly stop working out of nowhere. In most cases, there's a buildup of events leading to the actual outage, and that buildup is just as impactful as the actual incident to help determine the cause and the resources necessary to resolve the issue. A good platform will provide a timeline of events leading up to a major incident so that you're better prepared to prevent similar events in the future.

BEST PRACTICE

The ideal incident management platform should give responders full context of any incident, along with historical context that could help resolve an issue faster.

7 Does the platform provider have a community of DevOps users building atop their software with open source tools?

A platform provider that has an active community of users constantly expanding on top of it is one that is being improved on all the time. Having that community means your users can leverage the tribal knowledge of industry peers and leaders, and take advantage of tools for adjusting and customizing the solution to maximize ease and efficiency. The community also gives them a place to share and collaborate with others, and get feedback on their own tools and processes.
When choosing an incident management provider, look for one that:

- Is a proven thought leader with a track record of providing recommendations around incident response to ensure its customers' success
- Ensures that users have the opportunity to build on its platform and customize it via a well-documented API
- Is active in the community by attending DevOps meetups and can contribute meaningfully to a community of creative and talented engineers that are building open-source software

Encouraging the growth of this community involves creating a formal developer platform, which means going beyond creating a basic API with minimal documentation. More specifically, developers need development support, such as advanced documentation (including examples), self-service tooling, and dedicated service personnel.

The developer platform should provide examples and forums where community members can discuss API uses with each other to better streamline work and spur new ideas, creating a process for developers who create and maintain integrations. Downstream, customers need the ability to share, learn, and search for other users' work through a formal community environment—a single, vendor-neutral place for customers to discover, install, and read reviews of each app.

BEST PRACTICE

Look for a platform with an active user community that allows collaboration and encourages everyone to continually iterate on processes.

8 Can you gain actionable insights from data?

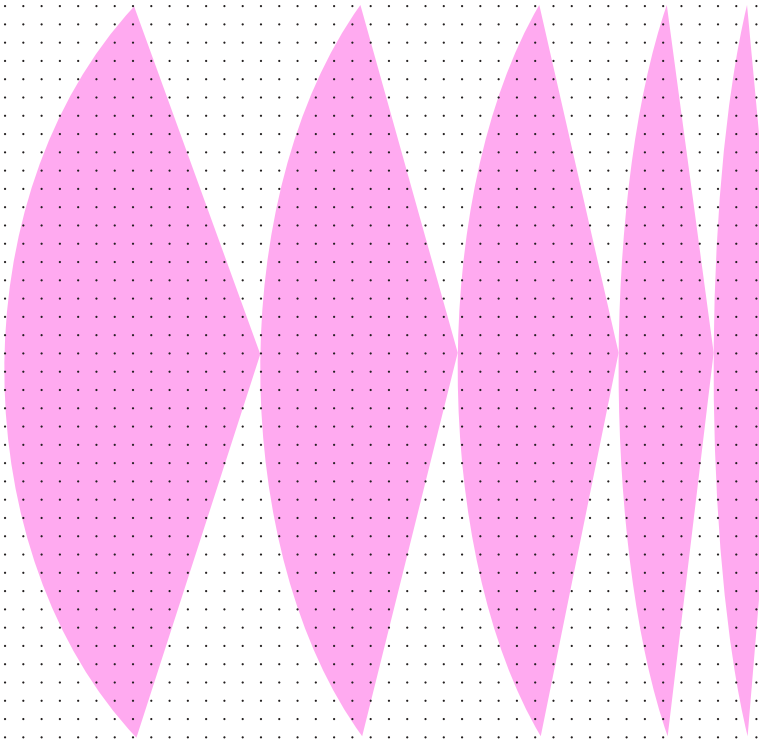

Organizations today have a lot of data they need to sort through in order to extract actionable insights like how to improve their operations or the health of their infrastructure. After all, what good is data if you can't do anything with it?

The platform you're considering should enable your team and organization to learn from peer and industry benchmark data. That means access to a large number of customers and a long history of collecting anonymized customer data. Not only that, it should spotlight areas for improvement.

The platform you're considering should ideally provide:

- **Scorecards to streamline the reporting process while pointing out key areas for improvement, such as alert fatigue and responder well being**
- **Service performance reviews for service execution to understand service performance and service risk moving forward. It also shows service correlation; for instance, does the shopping cart always goes down when you update the database?**
- **Business performance reviews to help business leaders review how operations affect business outcomes by presenting questions like:**
  - **Should you partner, build, or buy your next feature set?**
  - **What is the cost to maintain that new feature set?**
  - **Can your system handle more change or do you need to slow down your code update cycle?**
  - **Do you have enough resources to handle increased demand or do you have to increase headcount to handle more incidents?**

Choosing an incident management platform that's right for your team or organization is no small task, and it could be the difference between extended downtime and effective, fast incident response.

Curious to learn more? Visit pagerduty.com or sign up for a free, 14-day trial to see how PagerDuty can help you improve your incident response process.

PagerDuty, Inc. (NYSE:PD) is a leader in digital operations management. In an always-on world, organizations of all sizes trust PagerDuty to help them deliver a perfect digital experience to their customers, every time. Teams use PagerDuty to identify issues and opportunities in real time and bring together the right people to fix problems faster and prevent them in the future. Notable customers including GE, Vodafone, Box, and American Eagle Outfitters.
To learn more and try PagerDuty for free, visit www.pagerduty.com.
Follow our blog and connect with us on Twitter, LinkedIn, YouTube, and Facebook.

PagerDuty