

Modern Incident Management for IT Operations

A GUIDE TO OPTIMIZING IT OPERATIONS
AND DRIVING BUSINESS VALUE

Modern Incident Management for IT Operations

Abstract

For many employees or customers, IT operations is magic. Applications simply work, with little insight to the operations that actually allow systems to function. However, applications don't just magically work, especially in the complex, hybrid, IT environments of today.

This paper describes the challenges of incident management and provides advice to help you develop an operations competency to drive business value. Incidents will always happen, but how you handle them is the difference between a mediocre response and a great one.



Contents

4	Introduction
5	The Evolution of IT
6	Cloudy with a Chance of Downtime
8	The Not-So-Hidden Cost of Downtime
10	Didn't ITSM Solve Monitoring and Downtime Problems
12	The Evolution of IT Operations and Modern Incident Management
13	EBSCO Goes Agile
14	Modern Incident Management in Six Steps
17	The Opsgenie Advantage - The Incident Response Orchestration Platform
21	Incident Management Use Cases
22	Incident Management Success Stories
23	Conclusion
24	About Opsgenie

Introduction

There is no doubt that today's IT environment is more distributed, diverse, and flexible than ever before. Our workforce and customers are inherently mobile and 24x7 multi-tasking is the norm. It's not surprising to discover that the most successful companies are those that reliably fulfill millions of transactions in the blink of an eye, operate when it's convenient for the customer, and deliver new experiences at an ever-increasing pace.

This new business reality demands that IT must "go faster" and be more reliable than ever before, but traditional operational models, organizational structures, and alerting tools are struggling to adapt. Today's IT applications and infrastructure are very different from past solutions, and so too are the skills required to manage them. Outmoded practices often lead to tardy responses and poor business outcomes, which in turn lead to a reduction in customer loyalty or employee dissatisfaction.

Superior customer experience is the modern measure of success and great service depends upon the efficiency of incident resolution. As a result, incident management is now the essential competency for a thriving 21st century business.

“ Incident management is now the essential competency for a thriving 21st century business.



The Evolution of IT

We may think that subscribing to business applications is something new, but it's origins date back to the 1960s, and focused on mainframe or utility computing. Fast forward to just over ten years ago when the online retailer Amazon launched Amazon Web Services (AWS). AWS's goal was to offer IT infrastructure as web services, replacing up-front capital investments of server, network, and storage expenses with low variable costs that dynamically scaled as a business grew. You only paid for the computing resources that you actually used and nothing more.

Cloud computing of today has evolved beyond basic hosting, software as a service (SaaS), and infrastructure as a service (IaaS) offerings to become the engine of modern enterprise innovation. Developers are drawn to the cloud by the abundance of advanced services that can be incorporated into applications as diverse as machine learning to the internet-of-things (IoT).

Cloudy with a Chance of Downtime

The side effect of 40 years of computing innovation is that many companies now operate an eclectic mix of applications and systems. Some applications reside in their own data centers where they can be intimately controlled, whilst other applications are delivered on the cloud and management is delegated to the 3rd party application provider.


This collection of applications, services, and systems often results in a patchwork of solutions and processes for logging, monitoring, and alerting. It's not uncommon for an enterprise to use numerous monitoring tools to track thousands of application events or alerts per day. However, without a strategy for incident management, many are overwhelmed by the volume of events.


Many IT operations departments funnel the alerts into email boxes to counteract the volume problem, but are just making matters worse. These boxes need 24x7 monitoring by senior operations staff who can identify and escalate critical messages that need attention. One key challenge is that the process of monitoring email boxes doesn't scale. Staff need to sleep, take vacations, and miss work due to sickness. Monitoring email boxes also does little to help debug downtime-causing issues.

While infrastructure costs have declined, operations costs have risen - driven in part by the complexity of debugging issues when you don't control the entire system. Also, monitoring hybrid systems does little to assuage customer complaints when a system goes down. But at what cost does underdeveloped incident management mean to a company?

A Look at Overstock.com's Infrastructure

Overstock.com is a technology-based retail company that began business in 1999 as an online liquidator. Over the years, Overstock has grown far beyond those beginnings: expanding both its offerings and its business holdings. Indeed, as Overstock's business grew, so did its operational infrastructure. They used a vast array of both cloud and on-premises tools to ensure it could keep its business humming 24x7. The following list is just a sample of applications Overstock.com utilized to provide customers with unbeatable system reliability: Catchpoint, Consul, Email, Heartbeat, Icinga, Icinga2, Nagios, ServiceNow, Slack, Splunk, and various APIs for custom applications. Each application having its own administrative UI, alerting mechanisms, and system logs.





The Not-So-Hidden Cost of Downtime

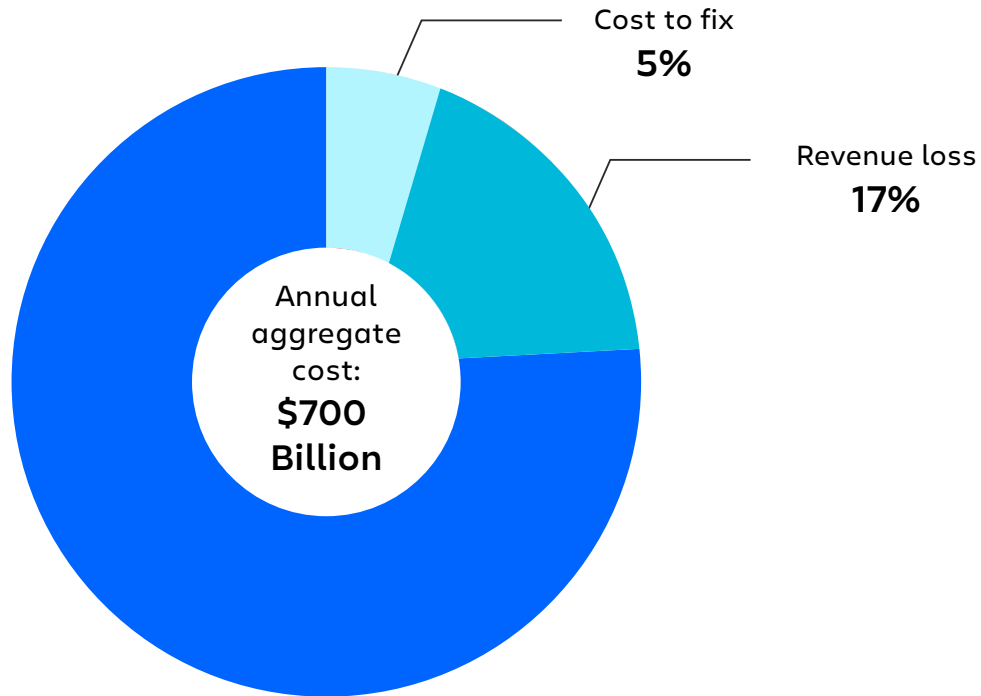
When outages occur they often affect the bottom line. The obvious fiscal damage is lost revenue, but to fully understand the cost of downtime we must also factor in other costs. There are operational costs to fix the outage by both the IT and business units, as well as the damage to reputation that could result in customer declines. Companies could also incur compliance and regulatory penalties.

Although there's no magic number on how much downtime impacts your business, Gartner estimates that the average enterprise incurs approximately \$5,600 for every minute of unplanned downtime in its primary computing environment. Gartner analyst, Andrew Lerner, further states that downtime at the low end can be as much as \$140,000 per hour, \$300,000 per hour in the middle, and as much as \$540,000 per hour at the higher end.

In another report IHS estimates that downtime costs North American organizations over \$700bn per year. IHS explored the frequency, length, cost, and causes of downtime and attributed 78% of the downtime cost to lost employee productivity. Lost revenue was attributed to 17% or \$119 billion. The actual cost to fix downtime was reported at 5% or \$35 billion.

IHS concluded that a typical mid-size company experiences 5 incidents and 27 hours of downtime per month, costing it an additional \$1 million per year. This cost balloons to over \$60 million a year for a large enterprise.

IT downtime costs North American businesses \$700 billion annually, mostly due to loss of employee productivity



©IHS, IHS Infonetics *The Cost of Server, Application, and Network Downtime: Annual North American Enterprise Survey and Calculator; 2016*

It's clear that there's a significant fiscal imperative to quickly find, diagnose, and repair IT incidents that cause outages and downtime.

Macy's Black Friday Nightmare

The Thanksgiving holiday is the most important period in a retailer's calendar. In 2017, according to the National Federation of Retailers, more than 174 million Americans shopped in stores and spent an average of \$335.47 per person online. However, Thanksgiving 2017 will be remembered by Macy's for a different reason. At around midday on Black Friday the retailer's credit card system appeared to struggle to process transactions causing outages on its website and long lines at stores. Six hours later the problem was resolved, but not before losing customers who abandoned their purchases and went on to other stores.

Didn't ITSM Solve Monitoring and Downtime Problems?

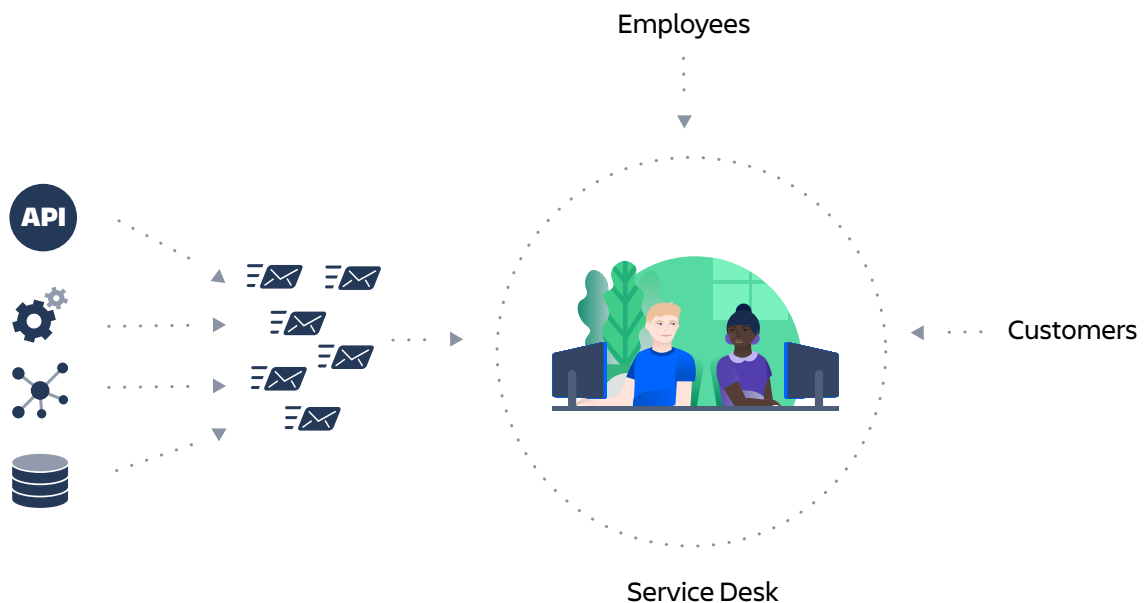
Historically, many operations teams regarded Information Technology Infrastructure Library (ITIL) as the single source of advice for managing their production environments and adopted a subset of the framework called IT service management (ITSM) to improve customer service. While it's true that ITSM is concerned with the implementation of IT services to meet customers' needs, it lacks any specific focus or recommendations for IT systems.

It's at this point that enterprises typically add an IT service desk. As defined by ITIL, the IT Service Desk is the primary IT function of ITSM. It is intended to provide a single point of contact for the communication needs of both users (i.e. customers), IT staff, and company employees. However, while adding a service desk manned by experienced operations personnel into the operations mix provides a central point of communication, it does little to help fix the cause of the downtime incident or address the following questions:

- What do you communicate about an incident?
- When should you communicate?
- With whom should you communicate?
- How frequently should you communicate?

The service desk adds extra pressure to an already overwhelmed operations staff because they are not only trying to fix unexpected outages but are also fielding service enquiries from personnel and customers. In this situation, it's easy for critical email alerts to be missed or simply overlooked. Also, multiple operations staff monitoring alerts from a single mailbox requires a high level of institutional cooperation and discipline to ensure that the same alert is not handled by more than one first responder.

Furthermore, service desk operations success is measured with metrics such as “call throughput” and “mean call time”. None of which contribute or directly measure the effectiveness of IT operations. Likewise, IT monitoring KPI's have little value if customer complaints increase.



Alert Fatigue

The never-ending stream of alerts that notify a team when there's a problem can be overwhelming. Because of this, team members responsible for handling alerts often experience longer response times, burn out, dissatisfaction with their work, and anxiety about failing to understand the real impact of the issue. Essentially, every alert may become meaningless because the monitoring software has “cried wolf” so many times.

The Evolution of IT Operations and Modern Incident Management

Until the last decade, responding to IT incidents was the primary job of operations teams. Organizations typically implemented a tiered team structure (i.e. Level 1, Level 2, Level 3) to respond to issues reported by customers or monitoring tools with the goal to minimize operational cost while maintaining service levels. Level 1 responders would be the most cost effective, entry-level employees and if they could not resolve an issue then they would escalate the response to Level 2. Level 2 were typically more senior and experienced employees, therefore more expensive resources. This escalation process would continue until the issue was usually resolved.

The rampant growth of “always-on services” developed with modern cloud architectures, containers, and microservices, has created a greater interdependency between systems. A slow response to any outage could immediately impact a company’s reputation as it is amplified via multiple social media channels. In addition, IT operations budgets were not increasing and staff were already operating at maximum productivity. It’s was also very evident that homegrown attempts to implement ITSM were overwhelmed by IT systems that produced gigabytes of alert data every day.

The combination of these external forces means the structure and tools of response teams has to change. The objective is shifting from purely cost effectiveness, to the timeliness of incident resolution. Time to respond and time to fix are now the organizing principles of operations teams.

Organizations are making changes to reduce the impact of downtime, from investing in early-detection capabilities to improving redundancy, training and hiring new people, and implementing incident response processes. IT operations teams are seeing productivity gains by using automation and artificial intelligence, thereby moving from reactive response to proactive cost-effective processes.

EBSCO Goes Agile



Organizations are making changes to reduce the impact of downtime, from investing in early-detection capabilities to improving redundancy, training and hiring new people, and implementing incident response processes. IT operations teams are seeing productivity gains by using automation and artificial intelligence, thereby moving from reactive response to proactive cost-effective processes.



Modern Incident Management in Six Steps

It's time to refocus your efforts by developing a strategic process around incident and management that reflects the new business realities of today. What follows are six focus areas to help reduce the likelihood of alert fatigue and improve your mean time to resolution (MTTR).

1. Identify Critical Systems and Consolidate Alerts

The primary culprit of alert fatigue is sending the same meaningless, non-actionable alert, over and over. Focus on the most critical systems, de-duplicate redundant notifications, and tackle the noisiest alerts first.

“ We were not using Opsgenie to begin with and teams were really struggling because the number of alerts were quite high. The decibel level of noise was tremendously high, not allowing teams to focus on the right alerts in a timely way. That's where Opsgenie did a fantastic job

Dipankar Biswass
Agile Development Director of Engineering,
EBSCO Publishing.

2. Scheduling Critical Resources and Use Appropriate Scheduling Model

It is crucial that after alerts have been identified and handled that there are sufficient experts available to take action. Therefore, it's important that any modern incident management platform has integrated scheduling capabilities. Ensure that the scheduling software can also support multiple scheduling models (like follow-the-sun) that match your organization's preferred mode of operation.

3. Use Automation – Filtering and Routing Incoming and Outgoing Alerts

Automation is the key to operational efficiency and productivity. If possible, you should automate every aspect of incident management from alert routing, notification deduplication, message workflows, conference bridge creation, status page updates etc. Automation is the crucial factor to reducing downtime.

4. Multi-channel Communication

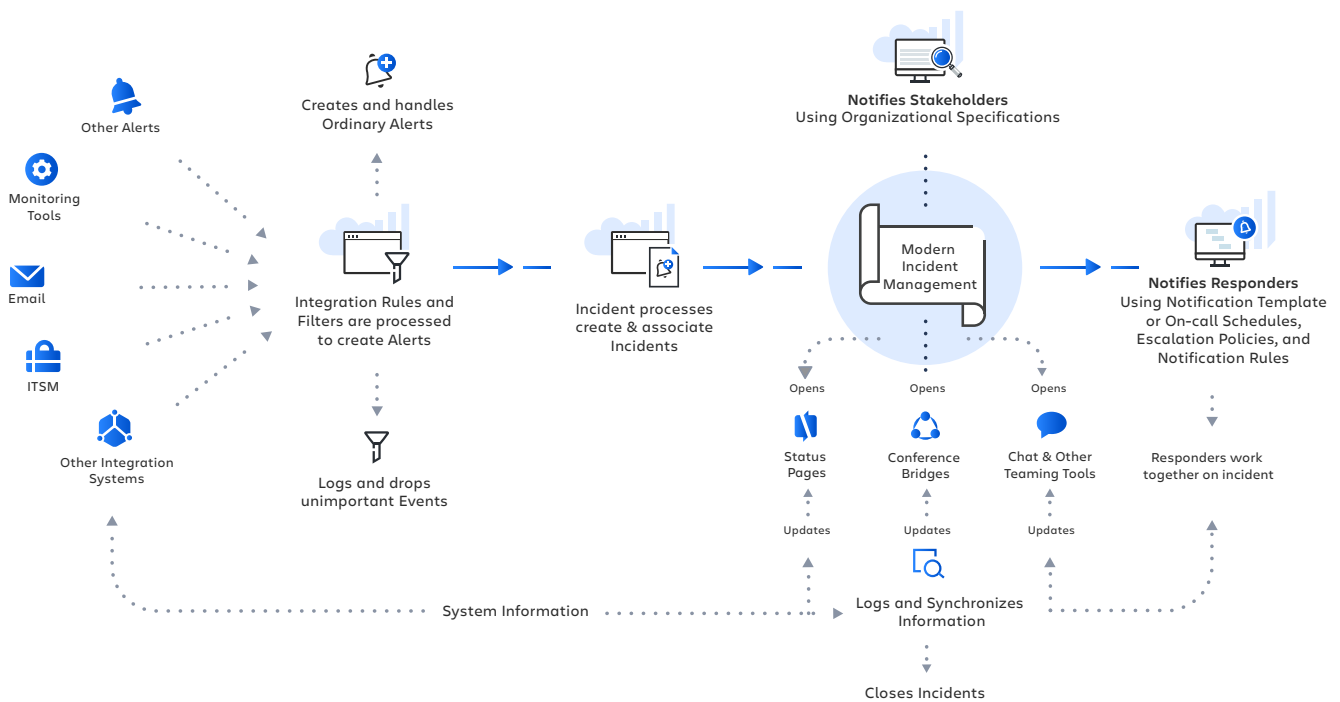
Remember that every incident affects three categories of constituents. The observers, responders, and stakeholders. Each one has a preferred communication channel that could include chat apps, emails, in-app notifications, telephone, and web pages. Also, you must reliably communicate with each constituent at the time of the incident, during the incident, and after the incident has been resolved. As mentioned above, it's helpful in automation to orchestrate the communication process.

5. System Monitoring

Efficiently responding to alerts and incidents is a core competency of incident management, so take time to monitor your systems performance. Implement a process whereby you can quickly view system health on a single pane of glass. Additionally, you should also automate receiving status reports on a regular schedule.

6. Post Mortem Analysis

After every incident or outage, expend extra effort to document what end user applications and or experiences were negatively influenced by the alert. Ensure that your incident management process has the capability to provide you with data to support post mortem analysis. This information will allow your support teams to improve their alert process and procedures, ultimately leading to an improved customer experience.



Companies that use a high degree of alert and incident automation are able to truly take advantage of the agility and innovation cloud computing provides and turn it into business value by deploying more features at lower costs of downtime per average deployment.

Responders – These are the members of your IT operations team that immediately respond to incidents. It's crucial for an efficient fix that the right individuals are notified in the correct order, using the most appropriate communication channel.

Stakeholders – These are members of your company, most likely in the management chain, who need to be kept informed about the incident status. The customer support manager for instance needs to know the status of the incident and the estimated time to fix. She can then prepare her customer support time to handle incoming customer calls.

Observers – These are your customers, employees, or end users who are ultimately impacted by the incident. Website visitors, for example, should be presented with information that explains the implication of the incident, rather than an unresponsive or broken website. Communicating with observers, even if it's "bad news", is crucial for delivering a superior customer experience.

The Opsgenie Advantage – The Incident Response Orchestration Platform

Rapidly identifying and responding to incidents in an increasingly complex infrastructure is critical to the survival of any organization. Time lost responding and resolving critical issues can result in dissatisfied customers and lost revenue. Opsgenie's Incident Response Orchestration platform helps organizations identify problems quickly and effectively, notify the right people, facilitate communications across business units, and collaborate to resolve problems fast.

Opsgenie is a cloud-based incident response platform with reliable, highly-available, distributed architecture replicated in multiple data centers and monitored around the clock. The life cycle for each alert, notification, and user action are recorded and reported to enable operations professionals to easily analyze IT incidents and outages. Opsgenie enables organizations to consolidate notification management into a single management system. No more time wasted digging through log files. No more finger pointing.

The Opsgenie Incident Response Orchestration platform contains the following innovative features that reduce your operational downtime:

- **Seamless Integration with Applications and Systems**

Integrate all of your applications and operations tools with ease. Opsgenie provides ready-to-use integrations with many operations services and tools. Opsgenie acts as the alert hub, receives alerts from monitoring tools and passes to team collaboration tools like Slack and HipChat.

Unbounce Chooses Opsgenie for Modern Incident Management

Unbounce was looking for alerting, on-call management, and escalations for distributed teams when they found Opsgenie in 2014. Previously, Unbounce's alerting and on-call management process was managed by multiple third party tools that were costly, and did not allow the team to customize the workflow to meet their needs. Mike Thorpe is Unbounce's Infrastructure Squad Manager. Speaking about Opsgenie's Slack integration Mike said, "We love the Slack integration with Opsgenie. We receive the ticket information, it then creates an event for our team as well as a list of items for us to review and take action. Tying both Slack and Opsgenie together was huge for our success as an organization."

- **On-call Scheduling Capabilities**

Easily create on-call schedules with daily, weekly, and custom rotations. On-call scheduling allows using multiple scheduling rules to use different coverage models at different times. Sophisticated scheduling scenarios such as after-hours coverage, weekdays, weekends, and geographical coverage for follow-the-sun can be defined quickly.

- **Response Orchestration**

Opsgenie orchestrates custom organizational response processes by notifying the right people based on alert policies and incident notification templates.

- **Response Collaboration**

Opsgenie initiates video and conference bridges automatically, as needed, for communicating about incident resolution. Notifications can also be routed to chat and other team collaboration tools.

- **Stakeholder Communication**

Opsgenie notifies stakeholders according to organizational specifications. Incident Status web pages can also be created automatically in order to notify stakeholders about incident resolution progress and service health.

- **Incoming Stakeholder Phone Call Routing**

Do you receive support calls from users during an outage? Incoming phone calls can be automatically routed to the right person using Opsgenie on-call schedules. If no one is available, Opsgenie takes a message, generates an alert, and notifies the correct responders.

- **Incident Reporting**

Opsgenie provides data visualization capabilities to give you a quick view into your organization's operational performance. There are two categories of Opsgenie reports:

GLOBAL REPORTS

Are account-wide and include generic analysis for notifications, API usage, and the overall alert responsiveness e.g. Current MTTA/R.

TEAM REPORTS

Focus on operations team activities, performance, and the alerts that they receive.

- **Incident Postmortem Analysis**

Opsgenie provides the tools to learn from past incidents, assess the efficacy of your incident response practices, and improve your team's future effectiveness.

IT Operations Metrics Glossary

MTTA – Mean Time to Assist

MTTF – Mean Time to Failure

FRT/MTTR – First Response Time/Mean Time to Respond

ATTR/MTTR – Average Time to Resolution/Mean Time to Resolution

= Total resolution time / Total number of requests

Incident Management Use Cases

Many of today's world class organizations have strategic initiatives that could be totally derailed if there are unplanned outages.

The following list describes some of those initiatives and where incident management plays a crucial role.

Digital Transformation and Infrastructure Modernization

Infrastructure modernization focuses on what is needed to support mission-critical and other applications and operations, addressing pressure driven by mobile and exploding data growth. In this scenario incident management is the crucial operational bridge between the old and the new. Operations teams are monitoring legacy application alerts while simultaneously providing support to brand new technology initiatives.

Customer Experience

Every retailer is looking to improve their customer experience, increase loyalty and maximize wallet share by using big-data to unify data and deliver real-time insights. Any downtime or drop in availability that affects this initiative not only interferes with data collection but could seriously impact the reliability of the eventual analysis. Incident management therefore plays a critical role in ensuring maximum uptime.

Cybersecurity

Modern day security threats are increasing in pace and sophistication, so many companies have deployed an SIEM system. Incident management is a perfect complement to these solutions that help you strengthen your security posture by efficiently automating the response to incidents generated by leading SIEM systems.

Internet of Things (IoT)

IoT has the potential to disrupt our homes, our commerce and every business, however any successful Internet of Things (IoT) initiative relies on the efficient communication between applications, devices, data and analytics 24x7. Downtime can't be tolerated. Incident management improves the operational efficiency of every IoT initiative through reduced downtime and increased device and system availability.

Incident Management Success Stories

POLITICO Runs Their Always-On Business with Opsgenie

POLITICO's systems need to be accessible 100% of the time to deliver quality news to their consumers. Before Opsgenie, POLITICO's internal help desk would receive alerts via email, perform a manual triage and assessment, then forward an alert email to the respective engineering team. Now that POLITICO has integrated with Opsgenie's Modern Incident Management platform, they have a much more efficient way to coordinate, administer and respond to alerts. POLITICO's team is also happy that they do not need to actively monitor a system when on-call since Opsgenie takes care of that for them.

Conclusion

There's no denying that the world of IT operations has undergone a massive shift within the past 10 years. Cloud computing has driven unprecedented flexibility, innovation and complexity into the enterprise, while increased customer expectations coupled with omnipresent social media have elevated customer service to the utmost priority. Today's IT applications and infrastructure are very different from past solutions, and so too are the skills required to manage them. Outmoded practices often lead to tardy responses and poor business outcomes, which in turn lead to a reduction in customer loyalty or employee dissatisfaction. As a result, incident management is the crucial discipline to implement to improve customer service and is the key differentiator of every 21st century business. Incidents will always happen but how you handle them is the difference between a mediocre business and a great one.

Next Steps

The World's top companies trust Opsgenie to support their IT Operations and DevOps teams because they know that Opsgenie provides the flexibility and power to securely and quickly respond to critical incidents that affect their service levels – and, more importantly, their customers. Take your incident management to the next level and try Opsgenie today, www.opsgenie.com.



About Opsgenie

Founded in 2012, Opsgenie is an Incident Response Orchestration Platform for IT Operations and DevOps teams. The solution controls all stages of incident response including alert management, intelligent on-call scheduling and escalation, and analytics for post-mortem incident analysis. With Opsgenie you can guarantee maximum uptime for your business by ensuring the right people are notified at the right time, with the most appropriate notification methods. Incidents happen. It's how you respond to them that matters.

