everbridge xmatters

Five Best Practices to Automate Major Incident Management

Using integrations and automation to reduce the impact of service disruptions

Introduction

How your company responds to major IT incidents makes a big difference in determining the extent of the business impact of these events. Managing these incidents effectively requires a well-coordinated set of processes, knowledgeable staff, and effective stakeholder communications. Even with these things in place, companies often struggle to execute in an organized way because they rely too heavily on manual activities.

According to a 2017 survey of DevOps organizations conducted by Atlassian and xMatters, companies report inconsistent processes, delays in declaring major incidents, delays in responding, and difficulty with effective collaboration. This echoed the results of a 2015 IT communications survey which found that more than 40% of businesses start to feel the impact of a major incident within 15 minutes after IT goes down, but 60% of companies cannot get the right person to respond to a notification that fast.

Major IT incidents take place within companies every day. While only a few make news headlines, IT outages, security breaches, and other incidents can cripple employee productivity, taint customer perceptions, and result in lost revenue.

You can set a goal of eliminating major incidents, but it's not really achievable. They will happen, and your best strategy is to be prepared for them when they do. Organizations that automate as much of their major incident resolution processes as possible achieve faster restoration of service and fewer mistakes.

A good starting place is understanding what a major incident is.

What are major incidents and why do they need to be treated differently?

After over 30 years of IT Service Management evolution, there is no generally accepted definition of what a major incident is or how to manage one. ITIL only gives major incidents a brief acknowledgment, telling companies:

- Agree what constitutes a major incident for your organization
- Maintain separate procedures for managing major incidents versus other incidents
- It is best to form a separate team to coordinate the major incident management process

So, what is a major incident then? Three commonly used criteria can help you make that assessment:

1. **Urgency:** The effect of the incident on deadlines and/or upcoming business events
2. **Impact:** The effects of the incident on business operations
3. **Severity:** The effects of the incident on users and customers

xMatters defines a major incident as one that impacts customers and their ability to do work.

What's the big deal about major incidents?

In a word, the big deal about major incidents is money.

According to the Ponemon Institute, the average cost of downtime in 2016 was $8,851 per minute. Do the math, and it comes to more than $500,000 per hour. With downtime averaging more than 90 minutes, you can see why a quick response is so important.

And that's just the immediate cost. The long-term costs can be unpredictable:

- Reputation damage
- Customer attrition
- Difficulty attracting new customers
- Fines and penalties (think GDPR)
- Customer remediation costs (especially in the event of a data breach)

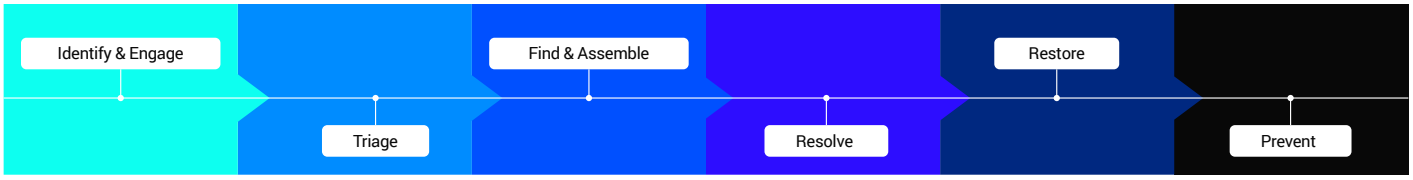

Basic Workflow for Major Incident Management

Most organizations' major incident management processes follow a similar basic workflow.

Identify & Engage: An incident is identified (by a user, helpdesk staff, or monitoring tool) and the Major Incident Manager is alerted to the potentially major issue.

Triage: The incident is assessed against major incident criteria and either accepted as a major incident or rejected as a 'false-alarm.'

Find & Assemble: The Major Incident Manager assembles a resolution team, engages stakeholders, and initiates major incident procedures.

Resolve: The incident is diagnosed, and corrective actions are implemented. This is where most of the difficult work takes place.

Restore: Service is restored to users and business operations return to normal.

Prevent: The Major Incident Manager and problem management teams conduct a post-incident review to identify preventative steps to avoid the incident from happening again.

Your primary goal should be shrinking the duration of the impact window, when your users and business operations feel the effects of the IT incident.

Major incident management issues

While the major incident management workflow seems to make sense, many companies report that the process isn't performing as well as they would like. The 2017 survey revealed:

- Process inconsistency – 50% have to wait for the operations center to declare a major incident.
- MTTR measurement mistakes – Most companies measure time to restore wrong by focusing on internal process time and SLAs instead of real impact time to users. The impact doesn't start when a major incident is declared by IT; it starts when the service goes down and includes monitoring, triage, and engagement time.
- Engagement bottlenecks – 34% say waiting for subject matter experts delays incident resolution.
- Duplicate work – 29% say duplicate tickets are created while an incident is being resolved, and tickets are routed without proper assignments and often must be rerouted.
- Communication is more important than the actual issue – Stakeholder perception is influenced more by the quality and effectiveness of communications than by the actual impact of the incident and time to resolve. Ironically, 43% of companies that report major incident communications are supported by manual processes.

A framework for optimizing MIM processes

The simplicity that gives the basic major incident management process its strength is also contributing to its weakness. Many companies are beginning to adopt a slightly more robust framework that includes three different timelines taking place in parallel as a foundation for their Major Incident Management optimization efforts.

Major Incidents Involving 3rd Party Suppliers

Incidents involving 3rd party service providers and external support teams require an additional layer of collaboration processes. Utilizing the same principles of planning ahead, negotiating SLAs in advance, setting expectations around resource engagement, and agreeing on communication roles and responsibilities can reduce confusion and conflict while a major incident is in progress.

Post-mortem reviews should be conducted with 3rd party suppliers to enable continuous improvement of the collaborative MIM process.

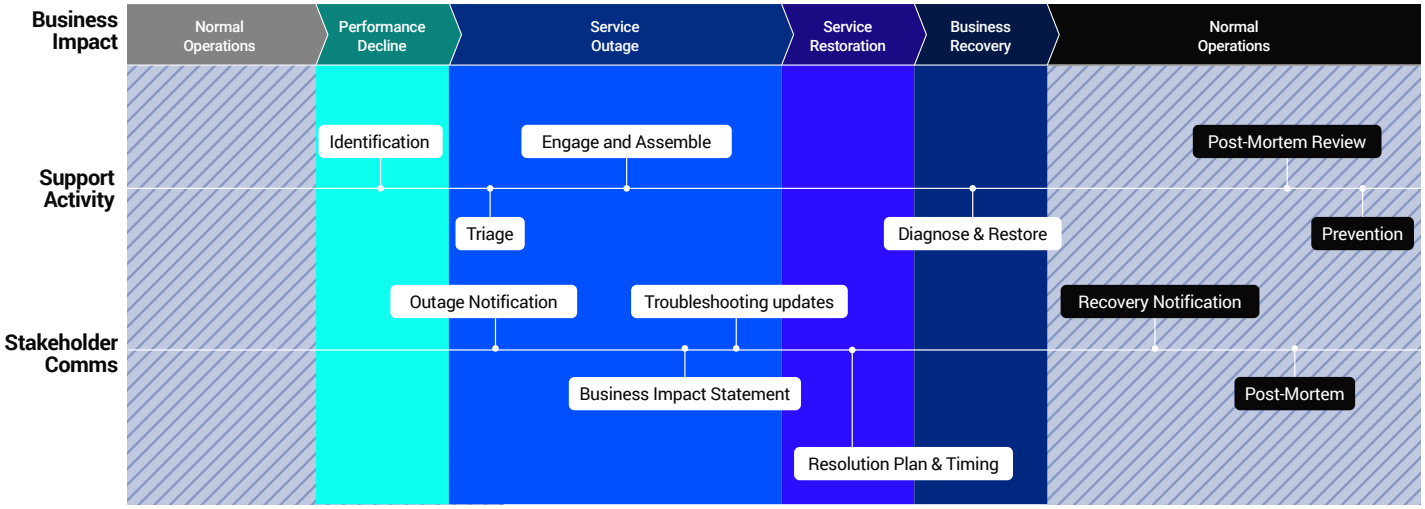

MTTR: Business Impact vs IT Engagement

5 best practices for moving activities outside the impact window

As you look to optimize the activities in your major incident management processes, the primary goal should be shrinking the duration of the business impact window - the period in which your users and business operations feel the effects of the IT incident. Look at what activities need to take place during the impact window and find a way to move other activities to either before the incident starts or after business has returned to normal operations. Here are some helpful suggestions:

1 Develop a process

Develop a major incident management process that encompasses what can be planned, coordinated, or executed during an incident. Identify support team contacts, including their skills and schedules, so your service desk can help them get started as quickly as possible. Set up your incident resolution team members with information about the incident so they can begin diagnosis immediately. Identify the stakeholders you will need to keep informed and develop message templates.

How to Automate:

- Sync your collaboration platform with your HR system and update regularly
- Correlate related alerts from your monitoring tool to automate context-rich service desk tickets and smart notifications
- Document the process somewhere accessible
- Authorize people to declare a major incident without waiting for a subject matter expert
- Run drills to practice the process

2 Get your infrastructure right

Apply filters to your monitoring alerts so your service desk technicians can more easily identify major incidents from routine noise. Collect data from systems and applications to support quick diagnosis and enable root cause analysis. Reduce the number of major incidents by acting when performance becomes degraded instead of waiting for the service to go down. Maintain a centralized issue tracking service so everyone involved can share information and operate from a single source of truth.

How to Automate:

- Apply filters to your monitoring alerts
- Use an APM solution that can crawl your applications and systems to identify the root cause
- Integrate your monitoring, service desk, collaboration, and chat solutions to share information
- Integrate your service desk, collaboration, and chat tools with your incident tracking solution

3 Measure time to restore service properly

Are you basing the time to restore (MTTR) on the period that IT is engaged or the period that the business is truly impacted? Measure the impact window from the business perspective to give the proper context for optimization efforts. Your goal is to minimize impact, not present better response reports to the board.

How to Automate:

- Provide visibility into a dashboard of applications to retroactively start the clock if necessary
- Preserve resolution activities and communications in the system of record for analysis

4 Keep stakeholders informed without interrupting issue resolution

Stakeholders expect effective and timely communications during major incidents – but they also expect your subject matter experts to remain focused on fixing the technical problems. You could consider designating a communications point of contact for your major incident management response team – someone who can monitor diagnosis and recovery activities, engage with business users to assess changes in impact, and prepare stakeholder communications without disrupting troubleshooting activities. A more effective strategy might be a web page with status updates so stakeholders could check for themselves without calling or emailing. Also, remember that stakeholder communication shouldn't stop when the service is restored – the most important stakeholder communication is the final one that summarizes what happened, what was learned, and how the situation will be prevented in the future.

How to Automate:

- Use a status update service to make incident status always available to stakeholders
- Build slash commands in your chat solution to update

5 Collect data to support problem management

Just because service is restored doesn't mean the support team can call it a day. Collect diagnostic and impact data to support problem management activities such as post-incident reviews to help avoid similar incidents in the future. By identifying what data needs to be collected to support problem management, the incident management team will not need to think about this during the impact period.

They can simply refer to a checklist and collect the data along the way – maintaining focus on their core objective of restoring service.

How to Automate:

- Preserve resolution activities and chat transcripts in the service desk or system of record for analysis
- Build a library of common occurrences and best practices for future events

It is important to be smart in how you connect your IT systems together to support critical processes like major incident management. Best practices for using technology to enable MIM include:

- Simplify - leverage consistent processes orchestrated with workflow automation
- Consolidate - create compound integrations amongst systems to enable a single process step to utilize multiple technology solutions
- Two-way notifications – enable both support staff and stakeholders to consume information and provide insights
- Use templates - ensure actionable alerts that are clear in their purpose and intent.
- Team empowerment - provide access to information through thoughtful automation
- Preserve collaboration records - use a system of record so transcripts can be used in the future to support problem management.

Leverage automation – how xMatters supports major incident management

Automation can play a powerful role in supporting major incident management processes. Automating everything might not be the best idea because major incidents can change in unexpected ways. Strike a balance between automation and efficiently curating manual work.

The xMatters service reliability platform uses integrations to enable your people work in their systems of choice. xMatters has more than 200 built-in integrations, plus an unlimited number of packaged and custom integrations. By enabling you to share information between tools, xMatters makes work more efficient, independent, and centrally controlled.

Identifying issues and engaging support

With xMatters, your monitoring and APM systems (like Splunk, New Relic, or AppDynamics) can share information about the offending event automatically with your service desk tool (like ServiceNow, Jira Service Desk, or Cherwell). This data sharing can help technicians more efficiently identify the issue and triage whether it is a major incident. If it is a major incident, automation can help service desk staff identify which individuals and teams need to be engaged for support and initiate the major incident process.

By syncing xMatters with your HR system, service desk technicians can see who's on call with the right skill set and preferred contact method, avoiding costly time engaging the right person. Integrations with service desk systems enable launching a conference bridge or chat room to targeted groups of people right from the service desk tool.

Enabling collaboration

Your support resources require an effective way to collaborate with each other. Instead of forcing an unfamiliar collaboration tool on them, let them use the tools they like! Many technicians prefer to use chat-based technologies to enable them to control the timing and flow of updates and provide an easy way of sharing diagnostic data with other troubleshooters.

With an integrated platform like xMatters, service desk technicians or engineers can open a chat room (like Slack or HipChat) for the event directly from whatever tool they're working in to get the conversation going easily and quickly.

In fact, people can use slash commands in Slack to open a conference bridge or save a chat transcript to the incident tracking tool (like Jira) so it can be used in a post-incident review. The transcript can help with improving processes and identifying preventative actions to avoid future issues.

Streamlining stakeholder communications

Many companies can manage status updates with customers and internal stakeholders by posting updates to a status web page (like StatusPage). With xMatters, they can automate this process and update directly through ChatOps commands or from the issue tracking tool. This can help ensure stakeholders get the timely updates they expect without causing interruptions in troubleshooting and resolving the technical issue.

Integrations with service desk tools can automatically include information from service desk tickets in communications to resolution team members, so they can avoid the need for copying and pasting key information (and sometimes get it wrong). As the incident progresses, new information can be captured in the issue tracking tool for centralized record-keeping and easy collaboration.

By integrating multiple systems into toolchains, xMatters enables full automation across the entire incident resolution lifecycle.

Major IT incidents take place within companies every day and can cripple employee productivity, taint customer perceptions, and result in lost revenue. But this doesn't have to be what happens to your company. Having a well-coordinated set of processes, knowledgeable staff, effective stakeholder communications, and the right set of technology capabilities from xMatters can turn a high-risk "disaster waiting to happen" situation into an opportunity for your IT staff and your company to shine.

About Everbridge

Everbridge, Inc. (NASDAQ: EVBG) empowers enterprises and government organizations to anticipate, mitigate, respond to, and recover stronger from critical events. In today's unpredictable world, resilient organizations minimize impact to people and operations, absorb stress, and return to productivity faster when deploying critical event management (CEM) technology. Everbridge digitizes organizational resilience by combining intelligent automation with the industry's most comprehensive risk data to Keep People Safe and Organizations Running™.

For more information, visit Everbridge.com, read the company blog, and follow us on LinkedIn and Twitter.

Get in touch to learn about Everbridge, empowering resilience.