fastly

Guide to incident command

Maarten Van Horenbeeck
VP of Security Engineering, Fastly

Anna MacLachlan
Content Marketing Manager, Fastly

Guide to incident command

Maarten Van Horenbeeck, VP of Security Engineering, Fastly
Maarten Van Horenbeeck (twitter: @maartenvhb) is Vice President of Security Engineering at Fastly. He is also Board member, and former Chairman, of the Forum of Incident Response and Security Teams (FIRST), the largest association of security teams, counting 300 members in over 70 countries.

Prior to his work at Fastly, Maarten managed the Threat Intelligence team at Amazon, and worked on the security teams at Google and Microsoft. Maarten has a master's degree in Information Security from Edith Cowan University, and a Masters degree in International Relations from the Freie Universitat Berlin. He enjoys backpacking, sailing and collecting first edition travel literature.

Anna MacLachlan, Content Marketing Manager, Fastly
Anna MacLachlan is Fastly's Content Marketing Manager, where she talks to brands and partners to tell stories about scale, security, and performance. She received her MA in Comparative Literature from NYU and loves megafauna and mountains.

© 2017 Fastly, Inc. All rights reserved

In any sufficiently complex system, failure is impossible to rule out, whether you're fighting literal fires or striving to keep your applications online. Failure can have a wide ranging impact, from undermining end users' trust to lost revenue. Fastly operates a large internetwork and a global application environment responsible for 1 terabit per second of internet traffic; as a result, we face many threats, both from a reliability and security perspective. As an organization, we've deliberately put in place a robust incident command system that allows us to rapidly identify and mitigate threats that come our way. We have these measures in place to ensure maximum team efficiency during incidents, with the overarching goal of working quickly and carefully to minimize risk and keep our customers online.

We also recognize the challenges that go with this, including communication, support, and proper delegating — making sure your sphere of responsibilities is taken care of while resisting the urge to jump in yourself, and empowering your teams to make smart decisions under pressure. We've designed our response to incidents to make the most of our team expertise, with the cornerstone belief that the people we hire are inherently intelligent, aware of their roles within the team, and are thus empowered to make decisions effectively.

In *Sources of Power*, Gary Klein studies groups from all walks of life — from firefighters and army generals to doctors and chess grandmasters — analyzing how they make critical decisions under pressure. While we'll explore his numerous findings in the next few chapters, one key takeaway is that we must trust our teams to make good decisions, having empowered them with the appropriate tools to do so. Instead of trying to construct fail-proof systems (they don't exist), we should trust the inherent competence of those systems' operators. Part of building that trust is making sure these operators "have the tools to maintain situation awareness throughout the incident"[1] — these tools include having an incident command structure to address threats when they do arise (as they will).

That said, you're not going to have a runbook for every scenario — you need to have a process for coming up with solutions for unexpected situations, and you can't be paralyzed when making a time-sensitive decision. Klein's studies of emergency responders showed that they don't compare various options against each other, but quickly evaluate courses of action by imagining how they might be carried out.[2] There is no rigid and careful analysis when fighting a fire — such measures would paralyze responders while they carefully weighed each option, wasting valuable time. Rather, firefighters evaluate incidents as they arise, opting for the best course of action without analyzing all available actions — "the emphasis is on being poised to act rather than being paralyzed until all the evaluations have been completed."[3]

There's no one way to approach incidents within your organization. Don't think of this text as a playbook to be followed rigorously, but rather as a metaplaybook: it's a way to think about how threats are identified and mitigated under pressure, without having to follow a step-by-step guide that would waste valuable time and resources. We put these structures in place at Fastly to empower our teams to work quickly to mitigate threats, not be paralyzed by the decision-making process. By providing the proper tools and training, we can fully trust our teams to make effective decisions as threats (like DDoS attacks) arise. Read on to learn how.

[1]Klein, 283
[2]Klein, 30
[3]Klein, 30

© 2017 Fastly, Inc. All rights reserved

Incident command

As an organization, we deliberately put in place a robust incident command system that would allow us to rapidly identify, mitigate, and contain incidents, and that would ensure effective communication flows both within the company and with our customers. In this article, we'll discuss the challenges a large global network faces, the protocols that we found helpful, and how you can apply them to your own organization.

Where to find inspiration

When you start developing a program, it's always good to look at how other types of people have solved similar problems. As engineers, we often tend to specialize and think the power to solve a particular problem is in our hands. When we do that, we can forget a few things: will we be able to ramp up engineers quickly enough, will our partners, such as network providers, be prepared and ready to help us when we need them to? How do we know if people on the frontline have the time and space to take care of basic needs (such as sleeping, eating, and de-stressing) during a prolonged incident?

It doesn't take very long before you realize established systems must already exist elsewhere: there's the Incident Command System (ICS) originally developed to address issues in the inter-agency response to California and Arizona wildfires, for instance, and the gold-silver-bronze command structure used by emergency services in the United Kingdom.[4] In addition, Mark Imbriaco's "Incident Response at Heroku" talk from Surge 2011[5] was a huge inspiration for our initial framework.

While technology has its own unique characteristics (like the ability to automate responses), many of the issues faced by these other responders still affect us today, such as communication between teams, making difficult decisions under pressure, and establishing hierarchy during an incident. There was no need to reinvent the wheel: we took some of the best practices of those who came before us when developing our own incident command system.

Understand what you're defending against

Another thing to understand well are the types of issues you're likely to face. Emergency responders determine the types of incidents they might encounter depending on their region and environmental factors — responders in Montana will face different issues in the summer (wildfires) versus winter (blizzards, power outages). We refer to our system as the Incident Response Framework (IRF), which came to be a catch-all for any issue that could directly cause customer impact, such as site downtime. Over time, as we professionalized, the system started specializing as well, and there are now specific plans in place covering smaller issues that may not yet cause customer impact, but may have the potential to do so in the future. In addition, a specific Security Incident Response Plan (SIRP) was developed that triggers on any issues which may have security repercussions.

We've engaged the IRF for security vulnerabilities that required immediate remediation, customer-visible outages, and issues of critical systems that support our network, such as our monitoring and instrumentation tooling.

When engaged, we identify the severity of an issue based on customer impact and business risk, as well as

[4] https://en.wikipedia.org/wiki/Gold–silver–bronze_command_structure
[5] http://files.meetup.com/2331301/incident_response_101.pdf

© 2017 Fastly, Inc. All rights reserved

the length of time the issue has manifested itself (see the severity matrix, below). Based on the severity, the team owning the affected service may be paged, or an organization-wide incident commander is allocated.

Identify the issue

Identifying an issue is critical — we have to know what we're dealing with. Within Fastly, we have multiple mechanisms in place to monitor service-related issues, including open source monitoring tools such as Ganglia, log aggregation tools such as Graylog and Elasticsearch, and several custom-built alerting and reporting tools. Ensuring events from each of these services make it to our service owners and incident commanders in a timely manner is critical so we can mitigate or avoid any customer impact.

SEVERITY	APP DELIVERY IMPACT	BUSINESS OPERATIONS IMPACT	SCOPE OF IMPACT
SEV0	Critical	Critical	All Sites Affected
SEV1	Critical	Critical	Multiple Sites Affected, or Single Site unavailable or suffering from severe degradation
SEV2	Major	Major	Multiple Sites Affected, Single Site intermittently available or suffering from minor degradation
SEV3	Minor	Minor	Single Site or limited customer impact

Every team owning a critical service at Fastly maintains a pager rotation, and receives reports regarding their own services directly. Engineering teams, however, are empowered to classify and re-classify events as needed to ensure pager rotations do not become too onerous. Most of them develop their own integration that ensures we don't suffer from alert fatigue on pageable events.

Events that do not lead to significant impact but could indicate wider problems over time are reviewed on a regular basis depending on their criticality, rather than leading to immediate action 24/7. As we covered last chapter, these can be recorded as part of modeling threats, for future exploration (and so they won't distract from the incident at hand).

Ramp up the right people

At Fastly, we hire and develop teams with care and intention — bringing the right group of people together is critical for well-thought-out decisions under pressure. We make roles and functions clear beforehand, and empower individuals to make use of their own expertise and trust their instincts — a combination which makes for efficient teams who make effective decisions.[6] And, it's important to remember the human element — we've seen attacks that have required all hands on deck, and it's critical that we don't burn people out. Engineers are humans too, and it's important to account for basic needs — such as eating and sleeping — when faced with an incident.

[6]Klein, 245

© 2017 Fastly, Inc. All rights reserved

Each team within Fastly designates a particular individual as being on call during a specific time slot. These people are primed to know that they will need to be more available and should keep their phones nearby. In addition, most teams that are critical for live services have a secondary engineer on-call, who is also aware of his or her responsibility to jump in in case of a major incident.

In addition, we maintain a few critical people on call for incidents that grow beyond a wider team or have customer impact. The role of these individuals is different — they know they won't be troubleshooting the issue directly, but they take on a number of critical roles that will help ensure mistakes are minimized, and investigation progresses as quickly as possible. They will:

- Coordinate actions across multiple responders
- Alert and update internal stakeholders, and update customers on our status — or designate a specific person to do so
- Evaluate the high-level issue and understand its impact
- Consult with team experts on necessary actions
- Call off or delay other activities that may impact resolution of the incident

Incident commander is not a role someone is ready to tackle when they're new to an organization. We select incident commanders based on their ability to understand Fastly's system architecture from the top down, so they have a mental model of how systems interrelate and impact each other — if you recall from last chapter, mental simulations and creating stories based on interrelated factors help decision makers envision the best course of action under pressure. Incident commanders are also well-versed in the structure of teams, so they know the right people to speak with and can reliably engage them. Finally, they're excellent communicators, and are able to maintain a cool, calm, and structured approach to incident response.

Above all, the incident commander is just that — a commander. They have the responsibility of ensuring the response is coordinated and running smoothly. To accomplish this, they must be empowered to delegate responsibilities and task people. They should be comfortable in this capacity and confident in their own abilities and their team. In *Sources of Power*, Gary Klein describes how experienced teams are well aware of the various roles and functions — and know the consequences when those break down.[7] Take for example the fireground commander who learned to bring a milk box with him to fires: when he first became a commander, he'd often leave his post to help put out the fire or offer medical attention. But, whenever he left his post his crew members couldn't find him when they needed a decision made, and wasted valuable time searching for him. Over time, his roles and responsibilities as a leader became clear to him, and he learned to keep a foot on the milk box to keep himself at his post. As Klein puts it, "He had realized the functions he served for his team and where he fits into the jobs of everyone else."[8]

A specific (positive) issue many incidents run into is volunteers — people who see a building on fire are often eager to help. The problem is, not everyone is equipped with the skill set for rescuing people, offering medical attention, or putting out fires. The same applies to your organization: when an incident is taking place, many of your employees will understandably want to help, even if they have no direct responsibility to being involved. When not properly managed, this can sometimes have negative effects: the environment can get overly chatty, or it's not clear who has picked up specific work. We've learned that removing people from the incident often is counterproductive — it demotivates people that want to work. Instead, we try to find opportunities to manage these volunteers, and either have them work on less critical items, or expand our scope of investigation beyond what we'd typically look at. This coordination will happen in a different room from the main incident, and is often coordinated by someone other than the main incident commander, but results are constantly communicated by a single individual.

[7]Klein 243
[8]Klein 243

© 2017 Fastly, Inc. All rights reserved

Communicate your status

Communication is critical, both in how we communicate incidents internally as well as to our customers. Poor communication — or worse, the lack thereof — leads to confusion and inefficiency, two things we can't afford when working quickly to assess incidents and keep our customers online. Both the method of how we communicate and what is communicated in these updates are important to consider — the tools you choose establish a framework for efficient incident response going forward, and what you communicate leads to fast, effective decision making.

Within Fastly, we use Slack and email as typical communication channels, and we use an external status-page hosted by Statuspage.io to communicate status updates to our customers to avoid any circular dependencies (i.e., if our ability to deliver web assets was in any way impacted, we'd still be able to communicate to customers). Our goal is to quickly publish status notifications to keep our customers informed — as with internal processes, poor communication leads to confusion. By keeping our customers in the know, we help inspire trust while we work to mitigate.

Interestingly, some of the services we rely on often are also Fastly customers. This means we can't necessarily depend on them being online during each type of incident affecting our service. As a result, we've grown through various backups, from our own IRC network, through phone chains, to alternative messaging tools to ensure systems are available. We also worked with some critical vendors to ensure our instance of their service operated on versions of their product that were not hosted behind Fastly to avoid these circular dependencies.

Over time we had to learn that various levels of individuals within the company have different needs for what they like to learn about an incident. In security incidents in particular, we assemble a specific group of executives who need to be briefed on very specific qualities of the incident — whether customer information was disclosed or not, or whether any critical systems were compromised.

Hence we've developed our processes to ensure incident commanders know what needs to be communicated, and to whom. During large incidents, quite often the incident commander will delegate ownership of communication to a dedicated resource to avoid over- or under-communicating an incident, which can erode the trust our customers place in us or lead to bad decisions.

Always improve

Each incident, as minor as it may seem, is logged in our incident management ticketing system. Within 24 hours after the incident, the incident commander will work with her or his team to develop an incident report, which is widely shared across the organization. We leverage the knowledge of the wider group involved to ensure it is as accurate as possible.

During this process, we use the time-proven "Five whys," a technique developed by Sakichi Toyoda of Toyota fame.[9] The idea is simple, and while there are no concrete rules, for every incident you ask why it took place, and for every answer you come up with, you ask the same question again. As you ask this question enough, usually about five times, you get to the actual root cause of the issue. The system is helpful in two ways — intermediate answers give us ideas about what we can do to mitigate a future incident, or monitor for it more effectively, but the last one tells us the underlying problem we likely need to address.

[9] https://en.wikipedia.org/wiki/5_Whys

© 2017 Fastly, Inc. All rights reserved

The root cause and each issue that hampered either the identification or response to the incident will receive its own ticket. Incidents that have unresolved tickets are tracked on a weekly basis in an incident review until all stakeholders are sufficiently assured that the right actions have been taken to prevent recurrence.

Incidents provide essential learning opportunities, often leading to new projects; brittle systems are often identified during these processes, and the additional visibility the organization gains often leads to the development of replacements or improvements.

© 2017 Fastly, Inc. All rights reserved