Incident Response in the Cloud Age

Version 1.3
Released: June 10, 2016

Securosis, L.L.C. 515 E. Carefree Highway Suite #766 Phoenix, AZ 85085 **T** 602-412-3051
info@securosis.com www.securosis.com

Author's Note

The content in this report was developed independently of any sponsors. It is based on material originally posted on the Securosis blog, but has been enhanced, reviewed, and professionally edited.

Special thanks to Chris Pepper for editing and content support.

This report is licensed by SS8.

www.ss8.com

SS8 is a time machine for breach detection. SS8 applies today's knowledge to history to find breaches now that you did not know about before. By generating, storing and analyzing months, and even years, of enriched intelligence from all communications flows, SS8 customers benefit from unprecedented content- and context-aware insights that allows them to find the threats that matter most. SS8 is trusted by six of the world's largest intelligence agencies, five of the 10 largest communications providers and two of the world's largest critical infrastructure entities. Learn more at www.ss8.com.

Copyright

Incident Response in the Cloud Age

Table of Contents

Shifting Foundations

Since we published our React Faster and Better research and Incident Response Fundamentals, quite a bit has changed for responding to incidents. First and foremost, incident response is a *thing* now. Not that mature security organizations didn't focus on responding to incidents before 2012, but since then a lot more resources and funding have shifted away from ineffective prevention towards detection and response. Which is awesome!

Of course, now that some organizations actually have decent response processes, the foundation is shifting beneath us. But that shouldn't be a surprise — if you want a static existence, technology isn't the best industry for you, and security is arguably its most dynamic speciality. We see the cloud revolution taking hold — promising to upend and disrupt almost every aspect of building, deploying and operating applications. Network speeds continue to increase, putting scaling pressure on every aspect of your security program, including response.

The advent of threat intelligence, as a means to get smarter and leverage the experiences of other organizations, is also having a dramatic impact on the security business — particularly incident response. Finally, the security industry faces an immense skills gap, which is far more acute in specialized areas… such as incident response. So whatever response process you roll out needs to leverage technological assistance — otherwise you have little chance of scaling to keep pace with accelerating attacks.

Not that mature security organizations didn't focus on responding to incidents before 2012, but since then a lot more resources and funding have shifted away from ineffective prevention towards detection and response. Which is awesome!

Entering the Cloud Age

There is this thing called the 'cloud,' which you may have heard of. As we have described for our own business, we are watching cloud computing change everything. That means existing I/R processes now need to factor in the cloud, which is impacting both architecture and visibility.

The cloud has two key impacts on your I/R process. The first is governance, as your data now resides in a variety of locations, with different service providers. You need various parties to participate to investigate an attack. The process integration of a multi-organization response is… challenging.

The other big difference is visibility, or lack thereof. You don't have access to network packets in an Infrastructure as a Service (IaaS) environment, nor can you see into a Platform as a Service (PaaS)

offering to see what happened. So you need to be much more creative about gathering telemetry on an ongoing basis, and figuring out how to access what you need during investigation.

Speed Kills

We have seen substantial increases in network speeds over the past 5 years, especially in data centers. So if network security monitoring is part of your I/R toolkit (as it should be), the way you architect your collection environment, and whether to actually capture and store full packets, are key decisions. Meanwhile data center virtualization is making it harder to know which servers are where, which makes investigation more challenging.

Getting Smarter via Threat Intelligence

Sharing attack data between organizations still feels a bit strange for long-time security professionals like us. The security industry resisted admitting that attacks succeed (ego got in the way), and feared (reasonably) that sharing company-specific data could provide adversaries with information to facilitate future attacks.

There just aren't enough skilled computer forensics specialists (who we call forensicators) to meet industry demand. You cannot just throw people at the problem, because they don't exist. So your team needs to work smarter and more efficiently.

The good news is that security folks got over their egos, a bit, and now understand that they cannot stand alone and expect to understand all the attacks that come at them every day. So sharing external threat data is now common, with both open source and commercial offerings available to provide insight, which is improving incident response. We have documented how the I/R process needs to change to leverage threat intelligence, so you can refer to that paper for detail.

Facing the Skills Gap

If incident response wasn't already complicated enough because of the changes described above, there just aren't enough skilled computer forensics specialists (who we call forensicators) to meet industry demand. You cannot just throw people at the problem, because they don't exist. So your team needs to work smarter and more efficiently. That means using technology more for gathering and analyzing data, structuring investigations, and automating what you can. We will dig into emerging technologies in detail later in this paper.

Evolving Incident Response

Like everything else in security, incident response is changing. The rest of this paper will discuss how. First we'll dig into impacts of the cloud, faster and virtualized networks, and threat intelligence on your incident response process. Then we'll discuss how to streamline a response process in light of the lack of people to perform the heavy lifting of incident response. Finally we will bring everything together with a scenario to illuminate the concepts.

More Data, No Data, or Both?

Given disruptions such as cloud computing and the availability of new data sources, including external threat intelligence, your incident response process necessarily needs to evolve. The following I/R process map shows how all the pieces fit together.

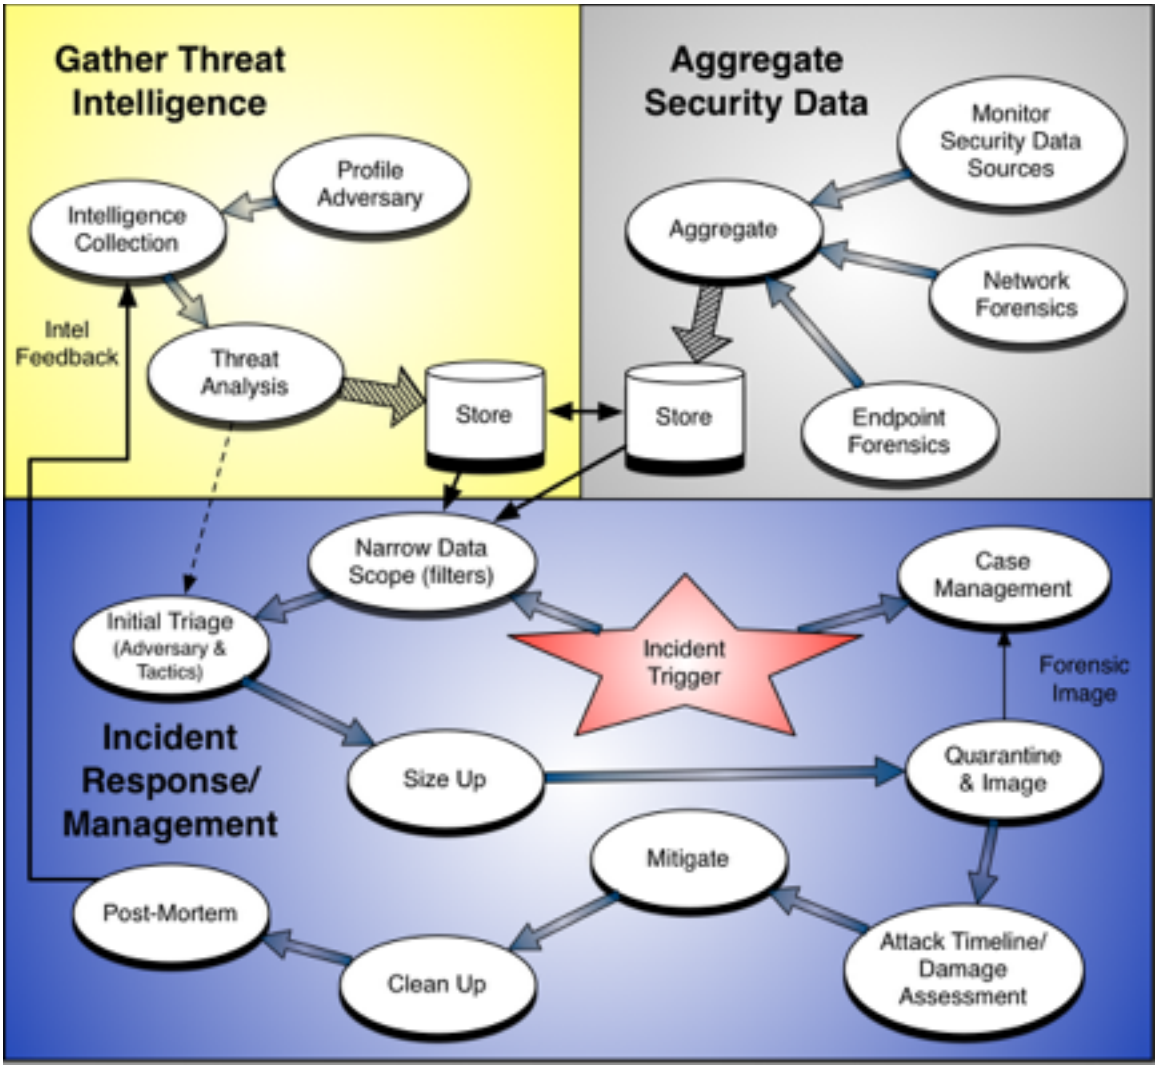

So what has changed in the two years since we last updated our I/R process map? Back then the cloud was nascent and we didn't know if DevOps was going to work. Today both the cloud and DevOps are widely acknowledged as the future of computing and how applications will be developed and deployed. Of course it will take time to complete this transition, but both are clearly real and in heavy use already, upending pretty much all the ways existing security works, including incident response.

Back then the cloud was nascent and we didn't know if DevOps was going to work. Today both the cloud and DevOps are widely acknowledged as the future of computing and how applications will be developed and deployed.

The good news is that our process map still shows how I/R can leverage additional data sources and functions to perform a complete and thorough investigation. It's hard to get sufficient staff to fill all the functions described on the map, but we will deal with that later in this paper. For now let's focus on integrating additional data sources, including external threat intelligence, and handling emerging cloud architectures.

More Data (Threat Intel)

We explained why threat intelligence matters for incident response in our TI+IR paper:

> To really respond faster you need to streamline investigations and make the most of your resources, a message we've been delivering for years. This starts with an understanding of what information would interest attackers. From there you can identify potential adversaries and gather threat intelligence to anticipate their targets and tactics. With that information you can protect yourself, monitor for indicators of compromise, and streamline your response when an attack is (inevitably) successful.

You need to figure out the right threat intelligence sources, then how to aggregate the data and run the analytics. We don't want to rehash everything in the TI+IR paper, but the most useful information sources include:

- **Compromised Devices:** This data source provides external notification that a device is acting suspiciously by communicating with known bad sites, or participating in botnet-like activities. Services are emerging to mine large volumes of Internet traffic to identify such devices.

- **Malware Indicators:** Malware analysis continues to mature rapidly, getting better and better at understanding exactly what malicious code does to devices. This enables you to define both technical and behavioral indicators, to search for across all platforms and devices within your environment, as described in gory detail in Malware Analysis Quant.

- **IP Reputation:** The most popular reputation data is based on IP addresses, and provides a dynamic list of known bad and/or suspicious addresses based on data such as spam sources, torrent usage, DDoS traffic indicators, and web attack origins. IP reputation has evolved since its introduction, and now features scores comparing the relative maliciousness of different addresses, factoring in additional context such as Tor nodes/anonymous proxies, geolocation, and device ID to further refine reputation.

- **Malicious Infrastructure:** One specialized type of reputation often packaged as a separate feed is intelligence on Command and Control (C&C) networks and other servers or sources of malicious activity. These feeds track global C&C traffic and pinpoint malware originators, botnet controllers, compromised proxies, and other IP addresses and sites to watch for as you monitor your environment.

- **Phishing Messages:** Most advanced attacks seem to start with simple email. Given the ubiquity of email and the ease of adding links to messages, attackers typically find email the path of least resistance to a foothold in your environment. Isolating and analyzing phishing email can yield valuable information about attackers and tactics.

As depicted in the map above, you integrate both external and internal security data sources, then analyze to isolate the root causes of the attacks, and figure out the damage and extent of compromise. Critical success factors in dealing with all this data are the ability to aggregate it somewhere, and then to perform the necessary analysis.

This aggregation happens at multiple layers of the I/R process, so you need to store and integrate all I/R relevant data. *Physical* integration is putting all your data into a single store, and then using it as a central repository for response. *Logical* integration uses valuable pieces of threat intelligence to search for issues within your environment, using separate systems for internal and external data. We are not religious about how you handle it, but there are advantages to centralizing all data in one place. So long as you can do your job, though — collecting TI and using it to focus investigation — either way works. Vendors providing big data security all want to serve as your physical aggregation point, but results are what matters — not where you store data.

Of course we are talking about a huge amount of data, so your choices for both data sources and an I/R aggregation platform are critical to building an effective response process.

No Data (Cloud)

What happens to response now that you don't control a lot of the data used by your corporate systems? The data may reside with a Software as a Service (SaaS) provider, or your application may be deployed in a cloud computing service. In data centers with traditional networks it's pretty

straightforward to run traffic through inspection points, capture data as needed, and then perform forensic investigation. In the cloud, not so much.

To be clear, moving computing to the cloud doesn't totally eliminate your ability to monitor and investigate your systems, but the insight traditional systems can provide into what's happening on those systems is dramatically limited.

To be clear, moving computing to the cloud doesn't totally eliminate your ability to monitor and investigate your systems, but the insight traditional systems can provide into what's happening on those systems is dramatically limited.

The first step for I/R in the cloud has nothing to do with technology. It's all about governance. Ugh. I know most security professionals just felt a wave of nausea. The G word is not what anyone wants to hear. But it is the only option for establishing rules of engagement with cloud service providers. What kinds of things need to be defined?

1. **SLAs:** One of the first things we teach in our cloud security classes is the need to have strong Service Level Agreements (SLAs) with cloud providers. And these SLAs need to be established *before* you sign. You don't have much leverage during negotiations, but you have *none* after you sign. These SLAs can include response time, access to specific data types, proactive alerts (them telling you when they had an issue), etc. We suggest you refer to the Cloud Security Alliance Guidance for specifics about proper governance structures for cloud computing.

2. **Hand-offs and Escalations:** At some point there will be an issue, and you'll need access to data the cloud provider has. How will that happen? The time to work through these issues is not while your cloud technology stack is crawling with attackers. Like all aspects of I/R, practice makes pretty good… there is no such thing as perfect. You need to practice data gathering and hand-off processes with your cloud providers. The escalation process within the service provider also needs to be very well defined to make sure you can get adequate response under duress.

Once the proper governance structure is in place, you need to figure out what data is available to you in the various cloud computing models and from your specific provides since that will vary. SaaS pretty much restricts you to logs (mostly activity, access, and identity) and information about access to the SaaS provider's APIs. This data is quite limited, but can help figure out whether an employee's account has been compromised, and what actions it performed. Depending on the nature of the attack and the agreement with your SaaS provider, you may also be able to get some internal telemetry, but don't count on it.

If you run your applications in an Infrastructure as a Service (IaaS) environment you will have access to logs (activity, access, and identity) of your cloud infrastructure activity at a granular level. A huge difference from SaaS is that you control the (virtual) servers and networks in your IaaS environment,

so you can instrument your application stacks to provide granular activity logging, and route network traffic through inspection/capture points to gather the packets for network security monitoring. Many IaaS providers also have fairly sophisticated offerings to provide configuration change data; some offer light security assessments to pinpoint potential security issues, both are useful during incident response.

Those running their own virtualization infrastructure (in either a private or hybrid cloud also have access to logs. As we have mentioned before, regardless of where the application runs, you can and should be instrumenting your applications to provide granular logging and activity monitoring to detect misuse. With limited visibility in the cloud, you don't really have much choice — you need to build security into your cloud technology stacks, and also ensure you are able to generate application logs to provide sufficient data to support investigation.

Capture the Flag

In the cloud — whether SaaS, IaaS, or hybrid — you are unlikely to get access to the full network packet stream. You will have access to specific instances (whether SaaS or hybrid), but obviously the type of telemetry you can gather will vary. So how much forensics information is enough?

- **Full Network Packet Capture:** Packets are useful for knowing exactly what happened, and being able to reconstruct and play back sessions. To capture packets you need either virtual taps to redirect network traffic to capture points, or to run network traffic through sensors in the cloud. But faster networks and less visibility make full packet capture less feasible.
- **Capture and Release:** This approach involves capturing the packet stream and deriving metadata about network traffic dynamics, as well as the content stream. It's more efficient because don't necessarily need to keep the full data stream, but you get much more information than you can glean from network flow data. This still requires inline sensors or virtual taps to capture traffic before releasing it.
- **Triggered Capture:** When a suspicious alert happens you may want to capture the traffic and logs before and after the alert on the devices and networks in question. That requires at least a capture and release approach (to get the data), with flexibility to only capture when you think something is important, so it's more efficient that full network packet capture.
- **Network Flows:** It will be increasingly common to get network flow data, which provides source and destination information for network traffic through your cloud environment, and enables you to see if there was some kind of anomalous activity prior to the compromise.
- **Instance Logs:** These are most similar to the output from the increasingly common endpoint detection and forensics offerings. If you deploy them within your cloud instances you can figure out what happened within the cloud instance, but might lack context on *who* and *why,* unless you also fully capture device activity. Understand that these tools will need to be work natively in the cloud, include support for autoscaling and virtual networking.

We have always been fans of more data, rather than less. But as we move into the Cloud Age practitioners need to be much more strategic and efficient about how and where to get the data to drive incident response. Some will come from external sources, and more from logical sensors and capture points within clouds, both public and private. The increasing speed of networks and broader and deeper telemetry available from instances (servers), especially in data centers, will continue to challenge the scale of data collection infrastructure, so scaling is a key consideration for I/R in the Cloud Age.

All this I/R data requires technology that can actually *analyze* it with reasonable timeliness. We hear a lot about "big data" for security monitoring these days. Regardless of what it's called by the industry hype machine, you need technologies to index, search through, and find patterns within data — even when you don't know exactly what you're looking for at the start. Fortunately other industries — including retail — have been analyzing data to discover and detect unknown patterns for years (they call it "business intelligence"), and many of their analytic techniques are usable for security.

We hear a lot about "big data" for security monitoring these days. Regardless of what it's called by the industry hype machine, you need technologies to index, search through, and find patterns within data — even when you don't know exactly what you're looking for at the start.

I/R collection in the cloud is more art than science, so you need to be constantly relearning how much data you need to capture to be able to do an adequate investigation. The feedback loop used to refine your I/R process is absolutely critical to make sure that when the right data is not captured, the process evolves to collect the necessary infrastructure telemetry, and/or instrument applications to ensure sufficient visibility.

Addressing the Skills Gap

Incident response in the Cloud Age requires evolution of the response process, to take advantage of data sources you didn't have before (including external threat intelligence), and new ways to analyze data that wasn't possible just a few years ago. You also need to factor in limited access to specific telemetry, particularly from the network, because you don't control the networks you're running on any more.

Lack of personnel is having a serious detrimental impact on pretty much every security organization. There simply are not enough skilled forensicators to meet demand.

But even with these advances, the security industry needs to face the intractable problem that comes up in pretty much every discussion we have with senior security types. Lack of personnel is having a serious detrimental impact on pretty much every security organization. There simply are not enough skilled forensicators to meet demand. And those who exist tend to hop from job to job, maximizing their earning potential. As they should, given free markets and all.

But this creates huge problems if you are running a security team and need to build and maintain a staff of analysts, hunters, and responders. Where can you find folks in a seller's market? You have a few choices:

1. **Develop them:** You certainly can take high-potential security professionals and teach them the art of incident response. Or given the skills gap, you may have to deal with other resources that may never become I/R ninjas. Sigh. This involves a significant investment in training, and a lot of the skills needed will be acquired in the crucible of an active incident.
2. **Buy them:** If you have neither time nor inclination to develop your own team of forensicators, you can get your checkbook out. You'll need to compete against consulting firms who can keep them highly utilized, which means they are willing to pay up for talent to keep the billable hours clicking along. And large enterprises can break their typical salary bands to get the talent they need as well. This approach is not cheap.
3. **Rent them:** Speaking of consulting firms, you can also find forensicators by entering an agreement with a firm that provides incident response services… which seems to be every security company nowadays. It's that free market again. This will obviously be the most expensive, because you are paying for the overhead of partners to bait-and-switch, and send you newly minted SANS-certified resources to deal with your incidents. That may be a little facetious, but only a bit.

You'll need all the above tactics to fully staff your team. *Developing* personnel is your best long-term option, but understand that some of those folks will inevitably head to greener pastures as soon as you train them up. If you need to stand up a team immediately, you will need to buy your way in and then grow. And it's a good idea to have a retainer in place with an external response firm to supplement your resources during significant incidents.

Changing the Game

It doesn't make much sense to play a game you know you aren't going to win. Finding enough specialized resources to sufficiently staff your team probably falls into that category. So you need to change the game. Thinking about incident response differently encompasses a lot, including:

- **Narrow focus:** As discussed earlier, you can leverage threat intelligence and security analytics to more effectively prioritize efforts when responding to incidents. Retrospectively searching for indicators of malicious activity and analyzing captured data to track anomalous activity enables you to focus efforts on devices or networks where you can be pretty sure adversaries are active.
- **On the job training:** In all likelihood your folks are not yet ready to perform very sophisticated malware analysis and response, so they will need to learn on the job. Be patient with your I/R n00bs and know they'll improve, likely pretty quickly. Mostly because they will have plenty of practice — incidents happen daily nowadays.
- **Streamline the process:** To do things differently you need to optimize your response processes as well. That means not fully doing some things you would, given more time and resources. You need to make sure your team doesn't get bogged down with things that aren't absolutely necessary, so they can triage and respond to as many incidents as possible.
- **Automate:** Finally you can (and will need to) automate the I/R process where possible. With advancing orchestration and integration options as applications move to the cloud, it is becoming more feasible to apply large doses of automation to remove a lot of the manual (and resource-intensive) activities from the hands of your valuable team members, letting machines do more heavy lifting.

Streamline and Automate

You can't do everything. You don't have enough time or people. Looking at the process map above, the top half is about gathering and aggregating information, which is largely not a human-dependent function. You can procure threat intelligence data and integrate it directly into your security monitoring platform, which is already collecting and aggregating internal security data.

Initial triage and sizing up incidents can be automated somewhat as well. We mentioned triggered capture, so when an alert triggers you can automatically start collecting data from potentially impacted devices and networks. This information can be packaged up and then searched for known

indicators of malicious activities or misuse (both internal and external), as well as compared to internal baselines.

At that point you can route the package of information to a responder, who can start to take action. The next step is to quarantine devices and take forensic images, which can be largely automated as well. As more and more infrastructure moves into the cloud, software-defined networks and infrastructure can automatically take devices in question out of the application flow and quarantine them. Forensic images can be taken automatically with an API call, and added to your investigation artifacts. If you don't have fully virtualized infrastructure, a number of automation and orchestration tools are appearing to provide an integration layer for these kinds of functions.

When it comes time for damage assessment, this can largely be streamlined due to new technologies as well. As mentioned above, retrospective searching allows you to search your environment for recognizable malware and behaviors associated with the incident under investigation. That will provide clues to the timeline and extent of compromise. Compare this to the olden days (like a year ago, ha!) when you had to wait for the device to start doing something suspicious, and hope the right folks were looking at the console when bad behavior started.

In a cloud-native environment, with an application built specifically to run in the cloud, there really isn't any mitigation or cleanup required, at least on the application stack. The instances removed from the application for investigation are replaced from known-good images which have not been compromised. The application remains up and unaffected by the attack. Attacks on endpoints still require either cleanup or reimaging, although endpoint isolation technologies make getting devices back up and running quicker and easier.

In terms of watching for the same attack moving forward, you can feed the indicators you found during investigation back into your security analytics engine to watch for them as things happen, rather than after the attack. Your detection capabilities should improve with each investigation, thanks to this positive feedback loop.

Magnify Impact

It makes sense to invest in an incident response management system/platform that structures activities in a way that standardizes your response process. These response workflows make sure the right stuff happens during every response, because the system requires it. Remember, you are dealing with folks who may be less experienced, so having a set of tasks for them to undertake, especially when dealing with an active adversary, can ensure a full and thorough investigation. This kind of structure and process automation can magnify the impact of limited staff with limited skills.

It may seem harsh, but successful I/R in the Cloud Age requires you to think differently. You need to make inexperienced responders more effective and efficient. On a scale of 1-10, you should look for people ranked 4-6. With training, a structured I/R process, and generous automation, you may be able to have then function at a level of 7-8, which is a huge improvement in effectiveness.

Cloud Age Incident Response in Action

Imagine you work for a mid-sized retailer which uses a mixture of in-house technology and SaaS, and has recently moved a key warehousing system to an IaaS provider as part of rebuilding the application for cloud computing. You have a modest security team of 10, which is not enough, but a bit more than many of your peers. Senior management understands why security is important (to a point) and gives you decent leeway, especially regarding the new IaaS application. In fact, you were consulted during the IaaS architecture phase and provided some guidance (with some help from your friends at Securosis) on building a resilient cloud network architecture, and how to secure the cloud control plane. You also had an opportunity to integrate some orchestration and automation technology into the new cloud technology stack.

The Trigger

You have your team on fairly high alert, because a number of your competitors have recently been targeted by an organized crime ring which has gained a foothold among your competitors; and proceeded to steal a ton of information about customers, pricing, and merchandising strategies. This isn't your first rodeo, so you know that when there is smoke there is usually fire, and you decide to task one of your more talented security admins with a little proactive *hunting*. Just to make sure nothing bad is going on.

The analyst got a hit on one of the samples, confirming your concern. You have an active adversary on your network. So now you need to engage your incident response process.

The admin starts poking around, searching internal security data with some of the more recent malware samples found in the attacks on the other retailers. The samples were provided by your industry's ISAC (Information Sharing and Analysis Center). The analyst got a hit on one of the samples, confirming your concern. You have an active adversary on your network. So now you need to engage your incident response process.

Job 1: Initial Triage

Once that you know there is a *situation,* you assemble the response team. You don't want to jump the gun until you know what you're dealing with, so you inform the senior team of the situation, but don't take any systems down. Yet.

The adversary is active on your internal network, so they most likely entered via phishing or another social engineering attack. Searches found indications of the malware on 5 devices, so you start capturing the network traffic from those devices. You want to be careful to avoid tipping off the adversary to their discovery.

Then you check your network security monitoring tool looking for indications that data has been leaking. You find a few anomalous file transfers, and luckily you integrated your firewall egress filtering capability with your network security monitoring tool. So once you detected the strange traffic being sent to known bad sites (via a threat intelligence integration on the firewall), you automatically started analyzing the network traffic from the devices which triggered the alert. Automation is sure easier than doing everything manually.

As part of your initial triage you got endpoint telemetry alerting you to issues, and network security monitoring data for a clue to what's leaking. This is enough to know you not only have an active adversary, but that more than likely you lost data. So you fire up your case management system to structure your investigation and store all the artifacts of your investigation.

Your team is tasked with specific responsibilities, and sent on their way to get things done. You make the trek to the executive floor to keep senior management updated on the incident.

Check the Cloud

The attack seems to have started on your internal network, but you don't want to take chances, and you need to make sure the new cloud-based application isn't at risk. A quick check of the cloud console shows strange activity on one of your instances. A device within the presentation layer of the cloud stack was flagged by your IaaS provider's monitoring system because there was an unauthorized change on that specific instance. It looks like the time you spent setting up that configuration monitoring service was well spent.

Security was involved in architecting the cloud stack, so you are in good shape. The application was built to be isolated. Even though it appears the presentation layer has been compromised, adversaries shouldn't be able to reach anything of value. And the clean-up has already happened.

Security was involved in architecting the cloud stack, so you are in good shape. The application was built to be isolated. Even though it appears the presentation layer has been compromised, adversaries shouldn't be able to reach anything of value. And the clean-up has already happened. Once the IaaS monitoring system threw an alert, that instance was taken offline and put into a special security group accessible only by investigators. A forensic server was spun up, and some additional analysis was performed. Orchestration and automation is a native feature of the cloud, and has huge value in accelerating incident response.

The presentation layer has large variances in how much traffic it needs to handle, so it was built using auto-scaling technology and immutable servers. Once the (potentially) compromised instance was removed from the group, another instance with a clean configuration was spun up to share

workload. But it's not clear whether this attack is related to the other incident, so you take the information about the cloud attack, pull it down, and feed it into your case management system. But the reality is that this attack, even if related, doesn't present a danger at this point, so you put it to the side while you focus on the internal attack and probable exfiltration on the internal network.

Building the Timeline

Now that you have completed initial triage, next you dig into the attack and start building a timeline of what happened. You start by looking at the comprised endpoints and network metadata to see what the adversaries did. From examining endpoint telemetry you deduced that Patient Zero was a contractor on the Human Resources (HR) team. This individual was tasked with looking at resumes submitted to the main HR account, and initial qualification screen for an open position. The resume was a compromised Word file using a pretty old Windows 7 attack. It turns out the contractor was using their own machine, which hadn't been patched and was vulnerable. You can't be *that* irritated with the contractor — it was their job to open those files. The malware rooted the device, connected up to a botnet, and then installed a Remote Access Trojan (RAT) to allow the adversary to take control of the device and start a systematic attack against the rest of your infrastructure.

You ponder how your organization's BYOD policy enables contractors to use their own machines. The operational process failure was in not inspecting the machine on connection to the network; you didn't make sure it was patched, or running an authorized configuration. That's something to scrutinize as part of the post-mortem.

Once the adversary had presence on your network, they proceeded to compromise another 4 devices ultimately ending up on both the CFO's and the VP of Merchandising's devices. Network forensic metadata shows how they moved laterally within the network, taking advantage of weak segmentation between internal networks. There are only so many hours in the day, and the focus had been on making sure the perimeter was strong and monitoring ingress traffic.

Once you know the CFO's and VP of Merchandising's devices were compromised, you can clearly see exfiltration in network metadata. A quick comparison of file sizes in data captured once the egress filter triggered shows that they probably got the latest quarterly board report, as well as a package of merchandising comps and plans for an exclusive launch with a very hot new fashion company. It was a bit of a surprise that the adversary didn't bother encrypting the stolen data, but evidently they bet that a mid-sized retailer wouldn't have sophisticated DLP or egress content filtering. Maybe they just didn't care whether anyone found what was exfiltrated, or perhaps they were in a hurry and wanted the data more than to remain undiscovered.

You pat yourself on the back, once, that your mature security program's egress filter triggered a full packet capture of outbound traffic from all the compromised devices. So you know exactly what was taken, when, and where it went. That will be useful later, when talking to law enforcement and possibly prosecuting at some point, but right now that's little consolation.

Cleaning up the Mess

Now that you have an incident timeline, it is time to clean up and return your environment to a good state. The first step is to clean up the affected machines. Executives are cranky because you decided to reimage their machines, but your adversary worked to maintain persistence on compromised devices in other attacks, so prudence demands you wipe them.

The information on this incident will need to be aggregated, and then packaged up for law enforcement and the general counsel, in preparation for the unavoidable public disclosure. You take another note that the team should consider using a case management system to track incident activity, provide a place to store case artifacts, and ensure proper chain of custody. Given your smaller team, that should help smooth your next incident response.

Finally, this incident was discovered by a savvy admin hunting across your networks. So to complete the active part of this investigation, you task the same admin with hunting back through the environment to make sure both this attack has been fully eradicated, and no similar attacks are in process. Given the size of your team, it's a significant choice to devote resources to hunting, but given the results, this is an activity you will need to perform on a monthly cadence.

Closing the Loop

To finalize this incident, you hold a post-mortem with the extended team, including representatives from the general counsel's office. The threat intelligence being used needs to be revisited and scrutinized, because the adversary connected to a botnet but wasn't detected. And the rules on your egress filters have been tightened because if the exfiltrated data had been encrypted, your response would have much more complicated. The post-mortem also provided a great opportunity to reinforce the importance of having security involved in application architecture, given how well the new IaaS application stood up under attack.

Another reminder that sometimes a skilled admin who can follow their instincts is the best defense. Tools in place helped accelerate response and root cause identification, and made remediation more effective. But Incident Response in the Cloud Age involves both people and technology, along with internal and external data, to ensure effective and efficient investigation and successful remediation.

Summary

To contain an advanced attack you need to Respond Faster and Better — detecting every attack before it happens is a pipe dream. By focusing on shortening the window between attack and detection (otherwise known as dwell time), and having a solid plan to contain and then remediate attacks, you give yourself the best chance to survive to fight another day. That is one of the most significant epiphanies security folks can have. You cannot win, so success is about minimizing damage. Yeah, that's crappy, but it is realistic.

Just when you thought you had your I/R process in decent shape, the technology infrastructure foundation started to shift dramatically. The inexorable move to the cloud is limiting visibility and access to compromised devices, adding the need to work with external service providers to get what you need. And that's all assuming you can find people to perform the investigation.

Yet, there are many reasons to be optimistic about the state of incident response. Threat Intelligence can make a big difference in understanding what to look for and the typical tactics of *your* likely adversaries to focus your response efforts. Combining the intelligence with advancing analytics helps you not just identify anomalous activity, but also an understanding of the mission and tactics of the adversaries. Increasingly sophisticated network security monitoring can provide tremendous insight into what happened during attacks, without necessarily having to keep the full packet stream.

And those who can embrace trustable automation are able to streamline their response efforts to make limited investigation teams far more efficient and effective. But those benefits require an institutional commitment to data collection at all levels of the computing stack, regardless of whether computation takes place within a traditional data center or the cloud.

But no collection of tools will ever replace a skilled team of incident handlers and investigators. Get the right people, establish the right processes, and then give them the tools and support to do what they do best.

If you have any questions on this topic, or want to discuss your situation specifically, feel free to send us a note at info@securosis.com.

About the Analyst

Mike Rothman, Analyst and President

Mike's bold perspectives and irreverent style are invaluable as companies determine effective strategies to grapple with the dynamic security threatscape. Mike specializes in the sexy aspects of security — such as protecting networks and endpoints, security management, and compliance. Mike is one of the most sought-after speakers and commentators in the security business, and brings a deep background in information security. After 20 years in and around security, he's one of the guys who "knows where the bodies are buried" in the space.

Starting his career as a programmer and networking consultant, Mike joined META Group in 1993 and spearheaded META's initial foray into information security research. Mike left META in 1998 to found SHYM Technology, a pioneer in the PKI software market, and then held executive roles at CipherTrust and TruSecure. After getting fed up with vendor life, Mike started Security Incite in 2006 to provide a voice of reason in an over-hyped yet underwhelming security industry. After taking a short detour as Senior VP, Strategy at eIQnetworks to chase shiny objects in security and compliance management, Mike joined Securosis with a rejuvenated cynicism about the state of security and what it takes to survive as a security professional.

Mike published The Pragmatic CSO <http://www.pragmaticcso.com/> in 2007 to introduce technically oriented security professionals to the nuances of what is required to be a senior security professional. He also possesses a very expensive engineering degree in Operations Research and Industrial Engineering from Cornell University. His folks are overjoyed that he uses literally zero percent of his education on a daily basis. He can be reached at mrothman (at) securosis (dot) com.

About Securosis

Securosis, LLC is an independent research and analysis firm dedicated to thought leadership, objectivity, and transparency. Our analysts have all held executive level positions and are dedicated to providing high-value, pragmatic advisory services. Our services include:

- **Primary research publishing**: We currently release the vast majority of our research for free through our blog, and archive it in our Research Library. Most of these research documents can be sponsored for distribution on an annual basis. All published materials and presentations meet our strict objectivity requirements and conform to our Totally Transparent Research policy.
- **Research products and strategic advisory services for end users**: Securosis will be introducing a line of research products and inquiry-based subscription services designed to assist end user organizations in accelerating project and program success. Additional advisory projects are also available, including product selection assistance, technology and architecture strategy, education, security management evaluations, and risk assessment.
- **Retainer services for vendors**: Although we will accept briefings from anyone, some vendors opt for a tighter, ongoing relationship. We offer a number of flexible retainer packages. Services available as part of a retainer package include market and product analysis and strategy, technology guidance, product evaluation, and merger and acquisition assessment. Even with paid clients, we maintain our strict objectivity and confidentiality requirements. More information on our retainer services (PDF) is available.
- **External speaking and editorial**: Securosis analysts frequently speak at industry events, give online presentations, and write and speak for a variety of publications and media.
- **Other expert services**: Securosis analysts are available for other services as well, including Strategic Advisory Days, Strategy Consulting engagements, and Investor Services. These tend to be customized to meet a client's particular requirements.

Our clients range from stealth startups to some of the best known technology vendors and end users. Clients include large financial institutions, institutional investors, mid-sized enterprises, and major security vendors.

Additionally, Securosis partners with security testing labs to provide unique product evaluations that combine in-depth technical analysis with high-level product, architecture, and market analysis. For more information about Securosis, visit our website: <http://securosis.com/>.