Incident response

A breach is in progress. Time for your incident response team to jump into action, but which way to jump? Evan Schuman explains.

You're sitting at your desk when the call comes. It might be from a payment card company telling you that your company appears to be the common point of purchase for a series of cyber attacks. Perhaps it's from the Secret Service, FBI or some other arm of law enforcement. They just finished searching a suspect's servers and they found a copy of your network diagram with network credentials. It might be a sinister-sounding voice wanting $10 million or your company's payroll will be published. It could even be one of your security contractors who hesitantly tells you, "This time, it looks like the real thing."

However you first learn that your company might have been breached doesn't matter much. You have a variety of actions to undertake right away and there's an excellent chance your incident response plan didn't anticipate your current situation.

Does it ever? A better name for those plans could be: "A list of very specific things to do in a situation very different from yours."

The overarching reality is that in the first hours and even days following athe detection of an incursion you truly know nothing. Were you perhaps breached more than a year ago and just learning of it now? Could someone on your team – intentionally or otherwise – be a factor? Not only do you know nothing in that first post-breach-discovery phase, but your initial probe might be more misleading than informative. So what should you do? Here are nine steps you could take to mitigate the situation.

1) The initial forensics report will be very wrong. One of the first things any CISO will want to know once you learn of a breach report is to have that initial forensics report done. But a common mistake is to take that initial report too seriously. That initial report – often completed within 48 to 72 hours of the team's activation – is frequently a thorough compilation of the evidence found on that initial log inspection, says Ed McAndrew, who in January 2016 stepped down as Assistant U.S. Attorney and cybercrime coordinator for the U.S. Justice Department.

Surrounding those kernels of truth are any deliberately deceptive clues left by the perpetrators, including deleted trails and manufactured red herrings. To take a physical crime scene analogy, step one is photographing the crime scene and step two is cataloguing what is there. Analyzing the results, looking for inconsistencies, comparing those records with backups and hundreds of other data points is the next, much longer phase.

It's not an indictment of the forensic team that the initial report is not correct, any more than it is criticism of a police photographer for capturing the bloody handkerchief the thieves planted at a crime scene to frame someone else.

But it does mean that typical behavior must be avoided. That means not using that preliminary report to brief your CEO, your

OUR EXPERTS: *Incident response*

Tim Cullen, a senior security consultant, Adapture
Mark Fidel, president, RiskSense
Darren Hayes, director of cyber security and an assistant professor, Pace University,
Ed McAndrew, private practices attorney; former Assistant U.S. Attorney and cybercrime coordinator for the U.S. Justice Department,
Alastair Paterson, CEO, Digital Shadows
Mark Rasch, former head of the U.S. Justice Department's cybercrimes group, currently a private practice attorney
Christopher Roach, national IT practice leader, CBIZ Risk & Advisory Services

125%

increase in zero-days in 2015 over the year before.

– Symantec, "Internet Security Threat Report," April 2016

board and certainly not the SEC or your customers. You know that you're going to have to walk back much of what you say, so why do it? Tell all of your constituencies the same thing, which happens to be the truth: We have received a report that we might have been breached and are currently investigating. We take such matters very seriously and we will update you as soon as we have a handle on what did and didn't happen.

Many companies will issue news releases – within a day of learning of the breach – that listed what the thieves did not access, as in: "The thieves did not access customer passwords." At that stage of the investigation, you might – at best – have a good sense of what in general was touched, but there's no way that you'll know with any certainty what wasn't touched. At this stage, that information is simply not yet knowable.

> "Most of the information you are going to get is misinformation."
>
> – Tim Callahan, SVP & global CISO, Aflac

Tim Callahan, the senior vice president and global CISO for Aflac, the $23 billion insurance giant, compares those initial post-breach-discovery hours and days to the fog of war. "It's a crisis situation, that initial moment when something happens and it's really foggy and you're trying to figure out what happened," he says. "Most of the information you are going to get is misinformation."

2) Whom do you initially trust? Nobody. Not even yourself. There is mixed advice on this point in terms of how you balance pragmatism and pure paranoia, which is a very common and useful trait in the immediate post-breach-discovery phase. On the maximum paranoia side, suggestions are that you choose three security forensics teams beforehand and that you'll randomly choose one of them to not just lead the initial probe, but to be the initial probe. In other words, this school of thought has you not using any of your salaried security lieutenants in the initial phase. Why?

There are three reasons why you would want to avoid using – in the initial phase – your salaried people or, for that matter, any contracted security people who handle routine daily security operations.

Let us consider the most paranoid position first. Given that you know nothing at this preliminary stage, you have to consider the possibility one of your people is directly involved, possibly as an accomplice or even a ringleader.

A semi-paranoid response is that one of your people might be unknowingly involved, such as if their network credentials were stolen. Your colleague might be innocent, but that initial report would look as though they were a perpetrator or, at best, reckless in protecting their credentials. Either way, it's a huge conflict of interest for that employee or contractor.

The most likely scenario is that someone cut some security corners. Employees are human and will, from time to time, do something against the rules that seems mild. For example, on a Friday night, an employee merely scanned some of the security logs for which they were responsible rather than reading them fully. As (bad) luck would have it, that report turned out to be the one where the first signs of the intrusion might have been detected. Would that have made a real difference? Hard to say, but that is another huge potential conflict of interest for that employee handling – or overseeing – a report potentially covering his or her own activities.

McAndrew advises admins to outsource immediately and keep your people out of it until you know more. "They have a bit of a conflict of interest," he says. "Information security personnel lose their jobs over some of these instances. This is personal. This isn't just business. This affects their careers, their standing in the organization."

However, Aflac CISO Callahan disagrees. "Without my internal team, I am not going

21%

Paper-based data security incidents handled by BakerHostetler law firm in 2014.

Source: BakerHostetler

to know," he says. "There simply has to be some internal analysis."

Isn't he worried about that conflict of interest? Not if the right controls are in place, particularly a privileged access solution that has the appropriate controls, proper alerts and proper audit capabilities.

Mark Rasch, an attorney specializing in security issues and the former head of the U.S. Justice Department's cybercrimes groupy, agrees with McAndrew that outsiders need to be used and, during those intense initial hours, used exclusively. "You don't put out your own fires," he says. "You call the fire department."

3) What to do when your cone of silence is broken. One of the most-cited law enforcement frustrations when they catch cyber thieves and get into their files is how often they find detailed and comprehensive information about a corporate victim's post-breach investigation, including email exchanges, text messages, recorded audio of phone conversations, copies of forensic reports and the like.

Here's a suggestion: Have a communication plan that is known in advance to the five or six people you plan on routinely communicating with post-breach-discovery. A critical part of that communications plan is to have throwaway phones, just like the bad guys do, McAndrew says. This would be a mobile device that is purchased solely for this purpose and is never used beforehand. All of the throwaways should be kept plugged-in inside a locked drawer.

Mark Rasch, private practice attorney

When the post-breach-discovery phase starts, all breach discussions either happen in-person (ideally in a room swept for electronic eavesdropping devices) or from one of these throwaway phones to another. Your office phones are potentially at-risk, especially if they use voice-over-IP (VoIP), McAndrew says.

And suspected networks are not solely at your office. This must be extended to any place where you have repeatedly worked before. If the bad guys tracked you in anticipation of the attack, they might have infiltrated those systems, too. That means Starbucks' network, your car's Wi-Fi, your home's network and even the network at your brother's house (if you've used it a few times before), McAndrew says.

4) Division of duties. Take a tip from NORAD. Some people can only order a launch and others can only execute a launch.

The internal team that is responsible for incident response should not have any administrative rights that would permit them conduct the kind of breach they investigate, says Terry Gold, founder and chief analyst at D6 Research (formerly IDanalyst). Once a breach is identified and the triage incident response team is dispatched, the CISO should bring in an external team of investigators to conduct the rest of the investigation.

"Audit should be involved to ensure that controls are in place and the processes are being followed," he says. While every company needs an incident response team that reports to the CISO, smaller organizations that have limited resources and staff will feel the pinch more because keeping the division of duties with a small staff can be problematic.

The IT staff that investigates potential breaches should not have the rights to access log files or the specific systems that store the data likely to be breached, he says. Even in a small company, it is a best practice to separate IT responsibilities, although some companies likely will end up blending the roles of watcher and those being watched.

While collusion is sometimes found in

30.6%

Companies in the UK that have incident response plans in place in 2015.

Source: Statista

insider threat scenarios, Gold notes that it is unlikely that it will take place among information security teams because these teams often are audited more than those in other departments.

5) Threats change behavior. Most of these tips speak to enterprise CISOs, but this one might speak to execs one or two levels up.

One of the more insidious recent trends among enterprise data breaches is the personal threat. This is where during a post-breach investigation the attackers learn the name of someone high up and use highly personal information to psychologically torment them – with the intent to make them ease up on the probe. Examples would be anonymous texts saying things like, "How's your Visa card ending in XXXX doing?" or "Noticed that your wife swiped her card at the Local Town Gym 30 minutes later than usual this morning. Is everything OK?" or "Seems that we just remotely accessed your security system and turned on all your lights. Sorry about that."

These mind games can seriously impact an executive's decision-making, especially if the executive is not accustomed to dealing with such direct and personal threats, McAndrew says. It might cause them to spend far more corporate funds on this defense than is fiscally warranted. Or it could just as easily have the opposite impact, where the target of the attack inappropriately underspends on the probe in an attempt to bend over backwards to not favor something that could be seen as defending one employee.

Terry Gold, founder and chief analyst, D6 Research

The security professionals' advice: If one of your direct reports has been so threatened, supervise them closely in a manner that you have never had to before. Ask them to explain and justify their decisions until the threat goes away.

6) How to disclose to customers? The glossy breach letter. Attorney Rasch recounts the tale of a breached retailer's marketing director. The marketing chief ordered that the letters to customers include a deep-discount-certificate and then the breach disclosure letter printed on glossy stock.

Why would a company announce a breach with a discount certificate and glossy promotion piece? "What do people do when they receive a coupon in the mail?" Rasch asked rhetorically. "They keep the coupon and throw the rest away, figuring that it's just marketing hype." The marketer was hoping to have breach-disclosure cake and to eat it, too. This way, she is in fact disclosing to customers, but most will never see it.

7) Follow the data. It's good protocol for companies to search for their data routinely and see where it might be hiding. Indeed, some companies have been known to seed their data with bogus information so that they can easily search for that data and they would know that it came from them. Or would they?

The security rationale for this tactic is to get an early warning of a breach by locating data remnants long before you could detect an intrusion. But just because the data originated on your site does not necessarily mean that your systems are the source of the leak. It might be a partner, such as a supplier, distributor, or mobile commerce vendor, who has access to your systems.

Alastair Paterson, chief executive officer of security vendor Digital Shadows, a San Francisco-based data analysis company, says that searches for data need to be extensive and must include dark websites as well as deep sites (not indexed by the search engines). "Oftentimes, clues are hiding in plain sight on public sites, such as Pastebin.

24

hours: The most important time after a breach to begin analysis and remediation

Source: Experian

Financial and personal data is also frequently sold via automated online shops, also known by international law enforcement as Automated Vending Carts," Paterson says.

But, sometimes, it's not a partner company at all but a wannabe partner. Paterson references a U.K. bank that found a folder with 3,000 of its internal documents floating on the internet, including the sensitive design of its ATM networks. The leak eventually was traced to some network-attached storage that was misplaced by an ATM firm that had merely bid on doing some ATM work for the bank. As part of the confidential bidding process, the bank shared technical details with bidders.

"It's not enough to just worry about your own systems," Paterson says.

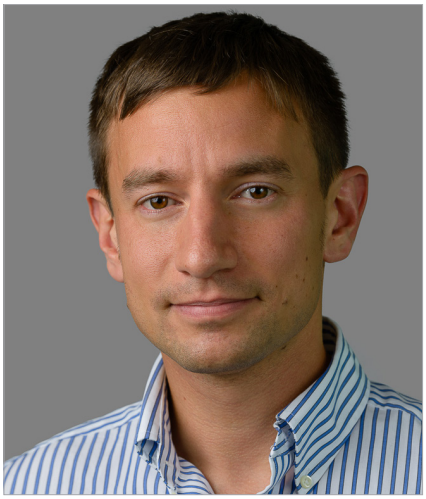

Alastair Paterson, CEO, Digital Shadows

A related tactic is pre-breach data-mapping. Although this does not directly protect existing systems, it can come in handy post-breach discovery. Think of it as the corporate equivalent of a homeowner doing a detailed list and video of everything in their house. In case of a burglary, that document and video file will deliver a much more comprehensive and emotionless report for police and the insurance company. After the breach, where can you turn for a secure version of what had been there? What if you're not sure if the backups were untouched?

8) In case of an ongoing attack, never ever shut down the network ...unless you have no choice. In the first hour after the attack is identified, stress will be high, especially if there is reason to believe the attack is ongoing. Darren Hayes, the director of cybersecurity and an assistant professor at Pace University, says it is critical to treat the servers as a crime scene before the police have arrived. "Don't let your security people touch systems and try and fix things," Hayes says. "Move everything to a backup site. Try and preserve as much as you can rather than deleting whatever you can. And that backup has to be isolated, it has to be off-network."

Tim Cullen, a senior security consultant at Adapture, a Norcross, Ga.-based IT consulting services provider, says his biggest fear is when IT does a kneejerk server shutdown. "Most people's immediate response to discovering a hack is to shut the server down. Your initial reaction is that you want to stop the bleeding immediately, but stopping that bleeding actually hurts your chances of recovery or prosecution," Cullen says. "You can lose or corrupt the file system to the point that forensics could be near impossible."

Even worse, attackers have been able to cut their own corners by planting a script that will delete all relevant security logs on reboot, Cullen says. In short, reboot before you've secured the system and the system might do its own sanitization.

Still, circumstances sometimes dictate that an immediate system shutdown is necessary, says Christopher Roach, the national IT practice leader for Cleveland-based CBIZ Risk & Advisory Services. "I'm not advocating that shutting down the server should never be done," he says. "A system shutdown is probably the most drastic responses to a security incident, but what if there were lives at risk? For example, a chemical plant gets a virus in its SCADA system. Letting the system run could lead to a major chemical disaster whereas shutting the system down or rebooting will trigger other safety measures that will avert a major chemical disaster."

9) Log file retention. How long is long enough? Although there are occasional instances of destructive cybercriminals breaking into a system to destroy or change files, most are infinitely more quiet. They watch activity, take notes and silently copy files from time

430M

new unique pieces of malware in 2015.

– Symantec, "Internet Security Threat Report," April 2016

to time. Indeed, if an attack is successful and professional, the victim may never know that he was attacked. Indeed, SMBs, which have no IT departments, may never know it. That means the investigators might need to review extensive log history – and far too many companies simply don't have it.

Federal prosecutor McAndrew says insufficient log files is the single biggest thing he'd want businesses to correct. "Usually, the instances have been going on much longer than originally known," McAndrew says. "We end up not being able to go back far enough to nail down exactly what happened. They're not saving enough evidence to really retrace the [attacker's] path. We need very fine granular data points. We sometimes see instances where there is no logging or [the triggers] are set so low that they are tripping out pretty quickly."

McAndrew recommends that companies retain logs for at least a year. Pace University's Hayes says he wants companies to go back far longer. "There should be no problem going back five years, given how cheap memory is today," Hayes says. ■

For more information about ebooks from SC Magazine, please contact Stephen Lawton, special projects editor, at stephen.lawton@haymarketmedia.com.

If your company is interested in sponsoring an ebook, please contact David Steifman, VP, publisher, at 646-638-6008, or via email at david.steifman@haymarketmedia.com.

1.5B

personal records were stolen or lost in 2015.

– Symantec, "Internet Security Threat Report," April 2016