Fact-Based Network Security: Metrics and the Pursuit of Prioritization

Version 1.2
Released: October 6, 2011

Securosis, L.L.C. 515 E. Carefree Blvd. 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.

Licensed by RedSeal Networks

RedSeal Networks develops proactive network security assessment software that enables organizations to visualize their security standing, maintain continuous compliance with regulations and better protect their critical assets. Unlike systems that detect attacks once they occur, RedSeal identifies holes in security infrastructure before they are discovered by hackers. RedSeal software analyzes and simplifies the complex interaction of firewalls and all other network security devices, delivering in-depth understanding of real-world exposure. For more information, visit RedSeal at www.redseal.net.

Contributors

The following individuals contributed significantly to this report through comments on the Securosis blog and follow-on review and conversations:

ds
Augusto Barros
Julia Allen
Russ Briggs

Copyright

Table of Contents

Introduction

Security programs at most businesses are about as mature as a pimply-faced teenager, which is problematic given the current state of security. Remember — attackers only have to get it right *once*, and some of them now hack more for Lulz than financial gain. How do you defend against an adversary who is more interested in pantsing you than stealing your stuff? But not all attackers fall into that category. You may also deal with state-sponsored adversaries — which means they have virtually unlimited resources while you don't. So you need to choose your activities wisely and optimize every bit of available resource just to stay in the same place. Unfortunately, far too many organizations don't choose wisely.

In fact, these organizations treat network security like Whack-a-Mole. Each time a mole pops above the surface, they try to it smack down. Usually that mole squeals loudest, regardless of its actual importance. But this means they spend a large chunk of time trying to satisfy certain people, hoping to get them to stop calling – and unfortunately that is much more about annoyance and persistence than actual importance of their demands. Sound familiar? Responding to the internal squeaky wheels clearly isn't a good enough prioritization scheme. Neither is the crystal ball, hocus pocus, or any other unscientific method. Clearly there must be a better way.

Responding to the internal squeaky wheels clearly isn't good enough. Neither is the crystal ball, hocus pocus, or any other unscientific method. **Clearly there must be a better way.**

Let's imagine a day when you could look at your list and immediately know which activities and tasks would cause the greatest risk reduction. How much would your blood pressure drop if you could tell the squeaky wheel that his top priority project was just not critical? With hard data to back it up? That's what Fact-Based Security is all about. Plenty of folks have metrics — but are they chosen and collected with an eye toward specific outcomes that matter to your business? You can get there by gathering metrics that guide and substantiate the decisions you need to make every day. Which change on which device is most important? Which attack path presents the biggest risk, and what's required to fix it? The data for this analysis exists, but most organizations don't use it. Shame on us, collectively.

In this paper we will investigate these issues and propose a philosophy to guide data-driven decisions. Of course we aren't talking about using SkyNet to determine what your security droids do on a daily basis. But your activities need to be weighed in terms of outcomes relevant to the business, which requires first understanding the risks you face – and more importantly assessing the relative value of the things you need to protect. Then we'll talk about what these reasonable outcomes should be, with the operational metrics to get there. Only once we have a handle on those issues can we talk about an operational process to underlie everything done with these metrics. With outcomes as a backdrop, using that data to make decisions can have a huge impact on both the effectiveness and efficiency of any security organization. Remember: *having* metrics and *using* them are totally different.

Then we'll dig into the compliance benefits of fact-based security, but for now suffice it to say that assessors love to see data – especially data relevant to good security outcomes. We'll wrap the paper by walking through a scenario where we actually apply these practices in a simple environment. That should give you the ammo you need to get started and to make a difference in your operational program(s). Let's go!

Defining "Risk"

Increasingly using data to determine your priorities enables you to focus on activities with the greatest business impact. But that begs the question: how do you determine what's important? The place to start is with your organization's assets.

Truth be told, both importance and beauty lie in the eye of the beholder, so this process challenges even the most clued-in security professionals. You will need to deal with subjectivity and the misery of building consensus (about what's important), and ultimately the answer will continue to evolve in light of the dynamic nature of business. But you still need to do it. You can't afford to spend a lot of time and money protecting devices no one cares about.

It's always good to start conversations with a good idea of the answer, so we recommend you start by defining relative asset value. We have long held that estimating (value = purchase price + some number you make up - depreciation) is ridiculous. That hasn't stopped many folks from doing it, but we'll just say there isn't much precision in that approach and leave it at that. So what to do? Let's get back to *relative*, which is the key.

We have long held that estimating (value = purchase price + some number you make up - depreciation) is ridiculous.

A reasonable approach would be to categorize assets into a handful of buckets (probably 3-4), distinguished by their importance to the business. For argument's sake we'll call them *critical, important, and not so important*. Start by spending time looking through your assets and sorting them into those categories. You can use a quick and dirty method of defining relative value which was first proposed in the Pragmatic CSO. Ask yourself and business leadership a few simple questions about the assets:

1. **What does it cost if this system goes down?** This is the key question, and it's very hard to get a precise answer, but try. Whether it's lost revenue, or brand impact, or customer satisfaction, or whatever – push executives to really help you understand what happens to the business if that system is not available. Is this information already available from a disaster recovery/business continuity analysis, or even a cyber-insurance underwriting exercise? Either might have produced some useful asset value data.

2. **Who uses this system and what data is on it?** This is linked to the first question, but can yield different and interesting perspectives. If five people in Accounting use the system, that's one thing. If every employee on the shop floor does, that's another. And if every customer uses the system and it

holds confidential data, that would be a much different thing. So a feel for the user community and stored data can give you an idea of a system's criticality.

3. **How easy are the assets to replace?** Of course having a system fail is a bad thing, but *how* bad depends on replacement cost. If your CRM system goes down, you can go online to something like Salesforce.com and be up and running in an hour or two. Obviously that doesn't include data migration, etc. But some systems are literally irreplaceable – or would require so much customization as to be effectively irreplaceable – and you need to know which are which.

Understand you need to abstract assets into something bigger. Your business leadership doesn't have an opinion about server #3254 in the data center. But if you discuss things like the order management system or the logistics system they can help figure out (or at least confirm) relative importance of assets. With answers to those questions, you should be able to dump each group of assets into an importance bucket.

You should be left with a basic understanding of your 'risk', which points out where to find the biggest steaming pile of security FAIL.

The next step involves evaluating the difficulty of attacking these critical assets. We do this to understand the negative side of the equation – asset value to the business is the positive side. If the asset has few security controls or resides in an area that is easy to access (such as Internet-facing servers), its criticality increases. So when we prioritize efforts, we can factor in both value to the business and likelihood of something bad happening if you don't address an issue.

And try to avoid self-delusion in this calculation. It's no secret that some parts of your infrastructure receive a lot of attention and protection while others don't. Be brutally honest about that, because it will enable you to focus on brittle areas as needed.

Like the asset side, focus on **relative** ease of attack and the associated threat models. You can use categories like: *Swiss cheese, home safe, bank vault, and Fort Knox*. And yes, we are joking about the category names.

You should be left with a basic understanding of your 'risk'. But don't confuse this concept of risk with an economic quantification, which is how most organizations define risk. Our risk assessment points out where to find the biggest steaming pile of security FAIL. This is helpful as you weigh the inflow of events, alerts, and change requests in terms of their importance to your organization.

And keep in mind that these mostly subjective assessments of value and ease of attack change frequently. That's why it's so important to keep things simple. If you need to go back and revisit the priority list every time you install a new server, it won't be useful for more than a day. So keep it high level, and plan to revisit these ratings every month or so.

Outcomes and Operational Data

Our next task is to determine the operational security metrics on which to base decisions, which requires as context the *outcomes* the business looks for. Outcomes are the issues central to business performance, and as such are both visible and important to senior management. Examples may include uptime/availability, incidents, disclosures, etc. Basically, outcomes are the end results of your efforts. You are trying to get to positive outcomes and stay away from negative outcomes.

These outcomes can be hard to decipher at times, if only because they tend toward the abstract. One way to stay grounded is to establish goals for improvement of these outcomes. This provides an idea of what you are trying to achieve and how to define success. To illustrate this, let's examine availability as an outcome – since it's rarely bad to improve availability of key business systems. Of course we are simplifying a bit – availability is more than just security. In the context of security, availability leads us to count issues/downtimes due to security problems.

Outcomes are the end results of your efforts. You are trying to get to positive outcomes and stay away from for negative outcomes.

Obviously many types of activities impact availability. Device configuration changes can cause downtime. So can unaddressed vulnerabilities that result in successful attacks. Application problems that may cause performance anomalies. Traffic spikes (perhaps resulting from a DDoS) can take down business systems. Even seemingly harmless changes to a routing table can open up an attack path from external networks and open the environment up to exploit. That's just scratching the surface. The good news is that you can leverage operational data to isolate the root causes of these issues.

What kinds of operational data do we need?

- **Configuration data:** Tracking configurations of network and security devices can yield important information about attack paths through your network and/or exploitable services running on devices.
- **Change information:** Understanding when changes and/or patches take place helps identify when devices need to be checked for proper/authorized changes and/or scanned again to ensure new issues have not been introduced.
- **Vulnerabilities:** Figuring out the soft spots of any device can yield valuable information about possible attacks.

- **Network traffic:** Keeping track of who is communicating with whom can help baseline an environment, which is important for detecting anomalous traffic and deciding whether it requires investigation.

Obviously as you go deeper into the data center, applications, and even endpoints, there is much more operational data that can be gathered and analyzed. But remember the goal. You need to answer the core question of "*what to do first*," establishing priorities among an infinite number of possible activities. We want to focus efforts on the activities that will yield the greatest favorable impact on security posture.

A simple structure for this comes from the Securosis Data Breach Triangle. In order to have a breach, you need data that someone wants, an exploit to compromise it, and an egress path to exfiltrate it. If you break any leg of the triangle you prevent a successful breach.

Data (Attack Path)

If the attacker can't see the data, they can't steal it, right? So you can focus efforts on blocking direct attack paths which would make it easy for attackers to access the data they want. Since you know your most critical business systems and their associated assets (from the risk definition task), you can watch to make sure attack paths don't open to expose this data.

How? Start with proper network segmentation to separate important data from unauthorized people, systems, and applications. Then constantly monitor your network and security devices to ensure there are no attack paths that put your systems and/or data at risk.

Given the complexity of most enterprise-class networks, this isn't something you can do manually, and it's most effective in a visual context. *Yes, in this case a picture is worth a million log records.*

Operational data, such as router and firewall configurations, is key for this analysis. You can also leverage network maps and ongoing discovery activities to check for the appearance of new attack paths. Any time there is a change to a firewall setting or a network device, revisit your attack path analysis. That way you ensure there's no ripple effect from a change that opens a window of exposure. Think of it as regression testing for network changes.

Given the complexity of most enterprise-class networks, this isn't something you can do manually, and it's most effective in a visual context. Yes, in this case a picture is worth a million log records. A class of analysis tools has emerged to address this. Some look at firewall and network configurations to build and display a topology of your network. These tools constantly discover new devices and keep the topology up to date. We also see evolution of automated penetration testing tools, which focus on continuously trying to find attack paths to critical data without requiring a human operator. There is no lack of technology to help model and track attack paths.

Regardless of the technology you select to analyze attack paths, this data helps determine what to fix first. If a direct attack path to important data results from a configuration change, you know what to do (roll it back!).

Likewise, if a rogue access point emerges on a critical network (with a direct path to important data), you need to get rid of it. These activities that make an impact and should take priority.

Exploit

Even if an attack path exists, it's not always practical to exploit the target. Here operational data on server configurations, as well as patch and vulnerability monitoring, pay dividends. Changes that happen outside authorized maintenance windows should arouse suspicion, especially on devices either containing or providing access to important data. Likewise, the presence of an exploitable critical vulnerability should bubble to the top of the priority list.

Again, if there is no attack path to the vulnerable device, clearly that lowers the priority of fixing the issue. So the prioritization decision needs to take into account data on all three legs of the Data Breach Triangle. Start with tracking what needs to be fixed on key business systems on a continuous basis to monitor their exploitability.

Egress Monitoring

In today's environment we cannot assume we know the attack vectors used by the bad guys. In fact, we shouldn't assume anything about anything. Attackers may harness a zero-day attack you haven't seen, which would evade a vulnerability and/or patch scan. Worse, a compromised device on your internal network renders attack path analysis somewhat irrelevant. The compromised device may already be in your house – which is where exfiltration comes into play.

At this point, you cease to deal with the theoretical (like attack paths) and need to deal with the cold, hard reality of your data leaving your network. You can do this through a few mechanisms – including network flow analysis, DLP, and content filtering. Or better: all of the above.

Monitoring network flows involves looking for strange source/destination pairs and sessions that violate typical traffic patterns. For example, it should make you a bit suspicious to see high traffic between a file server and an external device. Anomalous flows don't necessarily provide a smoking gun – it's hard to pinpoint what device the data actually comes from, especially in the face of a series of compromised devices during an exfiltration. But flows can certainly give you a good feel for what to investigate first.

At this point, you cease to deal with the theoretical (like attack paths) and need to deal with the cold, hard reality of your data leaving your network.

We also need to pay attention to content filtering on outbound devices, such as email and/or web security and DLP gateways. Here we are looking for examples of protected data leaving via the gateway, which generally indicates something bad. Remember, catching an issue on egress usually happens too late to actually stop the breach, but this analysis can shorten the window of exposure and give you a better opportunity to contain the damage.

Operationalizing the Facts

Do you want some good news? Most organizations do pretty well with the initial gathering of operational data. Early in the process, when the reports are new and the pie charts are shiny, it's easy to focus on collection and analysis of the data. Then the reality – the amount of work and commitment required to implement a consistent measurement and metrics process – sets in. Which is when most organizations lose interest and the metrics program falls by the wayside.

Then the reality – the amount of **work** and **commitment** required to implement a consistent measurement and metrics process – sets in.

Of course, given a clear and tangible connection between gathering data and doing your job better, you make the commitment and stick with it. That positive reinforcement can create a virtuous cycle. So it's critical, especially in the early phases of a fact-based network security process, to get a quick win and capitalize on that momentum to cement the organization's commitment to this model.

But consistency is only one part of implementing this fact-based network security process. In order to get a quick win and justify ongoing commitment, you need to make sense of the data. This issue has plagued technologies such as SIEM and Log Management for years: *having data does not mean you have useful and valuable information*.

We need to base decisions on facts, not faith. In order to do that, you need to make gathering security metrics an ongoing and repeatable process, and ensure you can interpret the data efficiently. This requires automation and visualization.

Automating Data Collection

Now that you know what kind of data you need, can you collect it? In most cases the answer is yes. From that epiphany, the focus turns to systematically collecting the types of data discussed above. Data sources such as device configuration, vulnerability, change information, and network traffic can be collected systematically in a leveraged fashion.

Inevitably the question of how deeply to collect data comes into play, whether you need to climb the proverbial stack in order to gather application and database events/logs/transactions/etc., or infrastructure level data suffices. In general, we Securosis folk advocate collecting more rather than less data. Not all of it may be useful now (or ever). But once you miss an opportunity to capture data you don't get it back. Of course which data sources to leverage depends on the problems you are trying to solve.

Remember, data does not equal information, and as much as we'd like you to capture everything, we know it's not feasible. So balance data breadth and fidelity against cost and storage realities. Only you can decide how much data you need to answer your questions and prioritize activities. We tend to see most organizations focus on network, security, and server configurations, logs, and events – at least initially. Mostly because that information is plentiful and largely useful in pinpointing attacks and substantiating controls.

Discussion of specifics for collecting and analyzing data of different platforms would be beyond the scope of this paper, but you should already know the answer is not Excel. There is just too much data to collect and parse. So at minimum you need to look for some kind of security management platform to automate this process.

Visualization

Next we come up against the seemingly intractable issue of making sense of the data you've collected. In practice, pinpointing anomalies and other suspicious areas which demand attention is much easier visually – so focusing on dashboards, charts, and reports become a key part of operationalizing metrics. It turns out those cool graphics available in most security management tools are more than eye candy. Who knew?

So which dashboards do you need? How many? What should they look like? It depends on which questions you are trying to answer. At the end of this paper, we will walk through a scenario to describe (at a high level, of course) the types of visualizations that are critical for detecting an issue, isolating its root cause, and figuring out how to remediate it.

But regardless of how you choose to visualize the data you collect, you need a process of constant iteration and improvement. It's that *commitment* thing again. In a dynamic world, things change constantly. That means your alerting thresholds, dashboards, and other decision-making tools must evolve accordingly. Don't say we didn't warn you.

In a dynamic world, things change constantly. That means your alerting thresholds, dashboards, and other decision-making tools must evolve accordingly.

Making Decisions

As we continue through our fact-based network security process, you now have a visual mechanism for pinpointing potential issues. But if your environment is like others we have seen, you'll have all sorts of options for what you can do. We come full circle, back to defining what is important to your organization.

Some tools have the ability to track asset value and show visuals based on them. Understand that value in this context is basically a totally subjective guess as to what something is worth. Someone could arbitrarily decide that a print server has equal importance to your general ledger system. Maybe it does, but this gets back to "relative value" from earlier in the paper. This understanding of an asset or business system's relative value yields key insight into how to prioritize your activities.

If the visualization shows something of significant value at risk, then fix it. Really. We know that sounds just too simple, and may even be so obvious it's insulting. We mean no offense, but most organizations have no idea what is important to them. They collect very little data and as such have little understanding of what is really exposed or potentially under attack. So they have no choice but to fly blind and address whatever issue comes next on the list, over and over again. Like Groundhog Day.

If the visualization shows something of significant value at risk, then fix it. Duh!

As we have discussed, that doesn't work out very well, so you need a commitment to collecting and then visualizing data, in order to even become aware of important issues. From relative value, you can prioritize activities by impact. You can do this systematically with automation to gain significant operational value.

But as one famous former CEO tends to say, there's "one more thing." You don't just get significant operational value from this automated data collection and visualization process. Having a fact base from which to make and prioritize decisions also goes a long ways toward substantiating the controls you have in place, making auditors happy. OK, maybe not *happy,* but less difficult during the audit process – which we'll discuss next.

Compliance Benefits

Beyond the operational value of fact-based network security, compliance efforts can benefit greatly from gathering data and being able to visualize and report on it. Why? Because compliance assessments force you to substantiate that your control set meets the spirit of whatever regulatory requirements you need to satisfy.

Let's run through a simple example. During a PCI assessment, the trusty assessor shows up with his/her chart of requirements. Requirement 1 reads "Install and maintain a firewall configuration to protect cardholder data." You have two choices at this point. The first is to simply tell auditor that you have and hope they believe you — not a recipe for success.

Or, you could consult your network security fact-base and pull a report on network topology, which shows your critical data stores (based on assessments of their relative value), the firewalls in place to protect them, and the flow of traffic through the network on the relevant attack paths to the critical assets/business systems.

Next the auditor needs to understand the configuration of the devices to make sure the firewalls block unauthorized protocols, protecting cardholder data. Luckily, the management system also captures firewall configurations on an ongoing basis. So you have current data on how the device is configured, and can show the firewall blocks the protocols in question. You can also explicitly show what IP addresses and/or devices can traverse it, using which protocols or applications (in the case of a new, fancy application-aware firewall).

The auditor may actually smile at this point, but will likely simply check the box in the chart and move on to the next requirement.

You close out this requirement by showing some event logs from the device, which demonstrate what was blocked by the firewall and why. The auditor may actually smile at this point, but will more likely simply check the box in the chart and move on to the next requirement.

Prior to implementing your fact-based network security process, you spent a few days updating the topology maps (damn Visio), massaging the configuration files to highlight the relevant configuration entries (using a high-tech highlighter) and finally going through a zillion log events to find a few examples to prove the implementation of the policies. Your tool doesn't make audit prep as easy as pressing a button, but it's a lot closer than working without tools.

Going where the money is

Compliance is a necessary evil in today's security world. Many of the projects we need to undertake have at least tangential compliance impact. Given the potential direct cost of failing an audit — having to disclose an issue to customers and/or shareholders, and fines — most large organizations have a pot of money to make compliance issues go away.

Compliance is a necessary evil in today's security world.

Smart security folks still think about Security First! Which means you continue to focus on implementing the right controls to protect the information that matters to you. But success hinges on your ability to show how the project can improve compliance, either by addressing audit deficiencies or making the compliance process more efficient to save money.

It's probably not a bad idea to keep time records detailing how long it takes your organization to prepare for a specific audit *without* automation. The numbers are likely to be shocking. In many cases, the real savings in time and perhaps resources can pay for the tools to implement a fact-based network security process. That would be nice, eh?

Fact-Based Network Security: In Action

Let's run through a simple scenario to illustrate the concepts we've discusses thus far. The key is simply to pick the most important item on the to-do list and do it. We make trade-offs every day. Some things get done, others don't. That's reality for everyone, so don't feel bad that you can't get everything done. Ever. But the difference between a successful security practitioner and someone looking for a job is consistently choosing the right things to get done.

Some folks intuitively know what's important and seem to focus on those things. They exist – we have met them. They are rock stars, but when you try to analyze what they do, there is no visible pattern. They are network security savants and just know. Sorry, but you probably aren't one of those folks. So you need a system – you know, a repeatable process – to make those decisions. You may not have finely tuned intuition, but you can overcome that by consistently and somewhat ruthlessly getting the most important things done.

Scenario: WidgetCo and the Persistent Attacker

In our little story, you work for a manufacturer and your company makes widgets. They are valuable widgets, and represent intellectual property that most nations of the world (friend and foe alike) would love to get their hands on. So you know your organization is a target.

The difference between a successful security practitioner, and someone looking for a job is consistently choosing the right things to get done.

Your management gets it – they have a well-segmented network, with firewalls blocking access to the perimeter and another series of enclaves protecting R&D and other sensitive areas. You have IPS on those sensitive segments, as well as some full packet capture gear. Yes, you have a SIEM as well, but you're currently revisiting that selection. That's another story for another day.

Your employees are reasonably sophisticated, but human. You run the security operations team, meaning that your folks do most of the management and configuration of security devices. Knowing you are a target means you need to assume attackers have compromised your network. But your tight egress filtering hasn't shown any significant exfiltration. Yet.

Your team's task list seems infinite. The myriad of requests to open and close firewall ports to support collaboration with specific business partners remain unaddressed. Your company's sales team needs access to a new logistical application so they can update customers on shipments of widgets. And of course, you use a certain flavor of two-factor authentication token to protect remote access for those reps.

Your boss lights up your phone almost daily because she gets a lot of pressure to support those business partners. Your VP of Engineering does some cool stuff with a pretty famous research institution in the Northeast. The sales guys are on-site and don't know what to tell the customer. And your egress filters just blocked an outbound attempt coming from the finance network, maybe due to the 2FA breach. What do you do? No one likes to be told no, but you can't get everything done. How do you choose?

No one likes to be told no, but you can't get everything done. How do you choose?

Get back to the risks

If you think back to how we define risk it becomes clear. Which assets are most important? Clearly it's the R&D information, which you know is the target of persistent attackers. Sure, you understand the value of customer information to them, and finance information would make some hedge fund manager another few million, but it would be bad if the designs for the next-generation widget ended up in the hands of a certain nation-state. Very bad.

When you think about the outcomes that are important to your business, protecting the company's IP is the first and highest priority. That's what the executive team told you. It supports your billion-dollar valuation, and senior management doesn't like to screw around with the valuation – given the amount of stock they collectively own. Thinking about the metrics that underlie various outcomes, you need to focus on indicators of compromise on those most sensitive networks. So gather configuration data and monitor the logs of those servers. Just to be sure (and to be ready if something goes south) you'll also capture traffic on those networks, so you can React Faster and Better if and when another alert fires.

It's also a good idea to pay attention to the network topology and monitor for potential exposures, usually opened by a faulty firewall change or some other change error. Your operational system gathers this data on an ongoing basis, so when alerts fire you can jump into action.

Saying No

In our scenario, the R&D networks are most critical, pure and simple. So you task your operations team to provide access to the research institution as the top priority. Of course, not full unfettered access, but access to a new enclave for research collaboration. After your team makes the changes, you do a regression analysis to make sure you didn't open up any holes, using your network security configuration management tool. No alerts fired and the report came back clean. So you are done at that point, right?

We don't think so. Given the importance of this network, you keep a subset of the ops team with their eyes on the monitors collecting server logs, IDS, and full packet capture data. You have also tightened the egress filters just in case. Of course folks get grumpy when the firewall blocks them, but you can't take any chances, not when your intellectual property is at risk. Without a baseline of the new traffic dynamics and a better feel for the log data, it's hard to tell normal from a problem.

Next up, you need to deal with the potential breach. So you install a full packet capture device on that network to start grabbing the traffic. You may not have internal expertise to do a full investigation (or those folks need to remain focused on the R&D network), so in that case we recommend you bring in a 3rd party forensics firm to do a quick investigation and analysis, because your team will be occupied for a couple days. This ensures a fundamental and widespread breach has not happened on the finance network.

Understandably, this decision makes the VP of Sales unhappy because his folks can't get access to the logistical information. They're forced to have a support team in HQ pull a report and email it to the reps' devices. It's horribly inefficient, as the VP keeps telling you. But the sales group will need to deal. They are at the bottom of the barrel in this case. that's not all. You also haven't been able to fully investigate the potential issue on the financial network, although you did install a full packet capture device on that network to start grabbing the traffic.

How do you justify these tradeoffs, especially to the grumpy VP of Sales? With data, or the lack thereof. During your risk analysis process, everyone agreed how bad it would be for the R&D network to be compromised, and a close second would be a breach of financial data (Wall Street doesn't like that too much). You'll need to remind folks on the senior team because they will squeal, but the priorities must hold. Unless they want to change your priorities, which happens from time to time – if the squealing is loud enough. But until you get different marching orders, stay true to your plan. Even if it means upsetting some heavy hitters.

Of course, if the forensics guys show that the data from the token vendor breach was used to gain a foothold in your finance department, you will redeploy some (if not all) of your operations folks to do a more thorough investigation, as *the breach then becomes the most clear and present danger*. This is a long way from whack-a-mole, eh?

You may be right. Or not.

Are these the right decisions or priorities? You won't know for months, if ever, and you are always open to second guessing because only hindsight is 20/20. But you have used your organization's priorities and operational data to prioritize decisions. You made the decisions without depending on intuition or faith, and sometimes you'll be right. Actually we think most of the time you'll be right.

Are these the right decisions? You won't know for months, if ever, and you are always open to second guessing because only hindsight is 20/20.

But not always. As long as you can defend your decisions with data, and show consistency in how you decide what to do based on priorities that the senior team has defined, they should support those decisions.

Remember: no one is right all the time. That's not a realistic definition of success. **In our opinion, the goal is to be consistent and make fact-based decisions, based on operational data reflecting the priorities of your organization.** In this paper, we have outlined a process to achieve this consistency, and shown the importance of systematically collecting and analyzing operational data. At the end of the day, you can only affect the things within your control – among which, how you allocate time for yourself and your team is an important lever, if not *the* most important. Because ultimately time is the only resource we can't make more of.

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

About the Analyst

Mike Rothman, Analyst/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 elQnetworks 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, L.L.C. 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 follow 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/or 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/>.