



Security Benchmarking: Going Beyond Metrics

Version 1.3

Released: June 16, 2011

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 nCircle



nCircle is the leading provider of automated security and compliance auditing solutions. More than 5,500 enterprises, government agencies and service providers around the world rely on nCircle's proactive solutions to manage and reduce security risk and

achieve compliance on their networks. nCircle has won numerous awards for growth, innovation, customer satisfaction and technology leadership. nCircle is headquartered in San Francisco, CA, with regional offices throughout the United States and in London and Toronto.

Copyright

This report is licensed under Creative Commons Attribution-Noncommercial-No Derivative Works 3.0.



<http://creativecommons.org/licenses/by-nc-nd/3.0/us/>

Table of Contents

Executive Summary	4
Security Metrics (from 40,000 feet)	6
Collecting Data Systematically	9
Sharing Data Safely	12
Defining Peer Groups and Analyzing Data	14
Communication Strategies	16
Continuous Improvement	19
Looking Inward	21
Benchmarking in Action	22
About the Analyst	25
About Securosis	26

Executive Summary

Relative to many information technology practices, security is immature. As such, there are few common practices of quantifying operational practices to evaluate performance. Beyond the overall immaturity of security practice, another contributing factor is the success criteria for security practitioners usually boils down to one simple question: “Are you compromised right now?” Everything else is noise — at least to most senior managers concerned about private data and intellectual property loss.

As an industry we need to start thinking about security as a practice, not a firefight. Quantification and comparison are critical steps on that journey.

That said, a key aspect of maturing our security programs must be the collection of security metrics and their use to improve operational processes. Even those with broad security metrics programs still have trouble communicating the relative effectiveness of their efforts — largely because they have no point of comparison. Thus when talking about the success/failure of any security program, without an objective reference point senior management has no idea if your results are good. Or bad.

Enter the Security Benchmark, which involves comparing your security metrics to a peer group of similar companies. If you can get a fairly broad set of consistent data (both quantitative and qualitative), then compare your

numbers to that dataset, you can get a feel for relative performance. Obviously this is very sensitive data, so due care must be exercised when sharing it, but the ability to transcend the current and arbitrary identification of problem areas as 'red' (bad), 'yellow' (not so bad), or 'green' (a bit better) enables us to finally have some clarity on the effectiveness of our security programs. Additionally, the metrics and benchmark data can be harnessed internally to provide objectives and illuminate trends to improve key security operations.

Let's discuss some typical security metrics and why they would be interesting to compare to others:

1. **Number of incidents:** Are you overly targeted? Or less effective at stopping attacks? The number of incidents doesn't tell the entire story, but how you fare relative to others is quite interesting.

2. **Downtime for security issues:** How effective you are at stopping attacks? And how severe is their impact? The downtime metric doesn't capture true business impact, but does get at the most visible impact of an attack.
3. **Number of activities:** By tracking activity at a high level, you can compare your team to other security teams to figure relative activity levels. With the increasing pressure on staffing, knowing your folks don't have a lot more to give can help make the case for adding headcount, or help you hang on to your stars.
4. **Budget efficiency:** Do you spend more or less money than other companies your size? Do you have more or less staff, as percentages of IT and total employees? What about your capital and operational spending? Obviously the finance team would be very interested in how you compare to peers financially. We understand this number would straddle the line between the internally-focused operational metrics and senior management-relevant metrics, but at the end of the day fiscal responsibility is critical for any overhead function, thus would be interesting to compare.

And this is just the tip of the iceberg. As you'll see, the ability to benchmark your environment opens up a world of possibilities for management and communication.

Benchmarking is not a new discipline, as it provides the main method for comparing financial performance. Company stock analysis is really just a big benchmark of commonly accepted metrics, such as profit and loss, income and expenses. Nor is benchmarking new in the IT space having been used for decades in the mainframe and networking worlds. But applying these comparison techniques to security is novel, and we have seen them used quite effectively when making a clear case for security investment. Security benchmarking is not a short-term fix — it's a long-term journey. One that requires commitment from senior management and an ongoing focus on applying lessons derived from data to refine operational activities, as well as a mechanism to push for accountability from all parts of the organization.

That said, we believe the time is now for organizations — both large and small — to start investing in security metrics and benchmarks. As an industry we need to start thinking about security as a practice, not a firefight. Quantification and comparison are critical steps on that journey.

This paper will make the case for organizations of all shapes and sizes to start down the security benchmarking path. First we'll examine the underlying security metrics, discuss how to determine an appropriate peer group, and then delve into how the benchmark can be used to improve the operational excellence of any security program.

Security Metrics

(from 40,000 feet)

Any metrics program has multiple consumers so you need to package and present the data to these different constituencies. As you'll see, there is no lack of things to count. But the fact that you *can* count something doesn't mean you *should*. So let's dig a bit into what you can count.

Disclaimers: we can only go so deep in a white paper. If you are intent on building a security metrics program, you **must** read Andy Jaquith's seminal work [Security Metrics: Replacing Fear, Uncertainty and Doubt](#). The book goes into great detail about how to build a security metrics program. The first significant takeaway is how to define a good security metric in the first place. A good security metric must be:

1. Expressed as numbers
2. Have one or more units of measure
3. Measured in a consistent and objective way
4. Gathered cheaply
5. Contextually relevant

Contextual relevance tends to be the hard thing. As Andy says in [his March 2010 security metrics article in Information Security magazine](#): *"the metrics must help someone—usually the boss—make a decision about an important security or business issue."* That's where most security folks tend to fall down, focusing on things that don't matter, or drawing suspect conclusions from operational data. For example, generating a security posture rating from AV coverage won't work well. Alternatively, looking at defining a set of business oriented security metrics and finding a reasonable base for comparison can greatly ease the difficulty in making a case to introduce new technologies, processes or procedures into the organization. To be clear, the security metric needs to relate back to a business imperative, perhaps keep business systems available. Thus security metrics which relate to ensuring those business systems remain available can substantiate the need for the technology.

Consensus Metrics

We also need to tip our hats to the folks at the Center for Internet Security, who have published [a good set of starter security metrics](#) built via their consensus approach. Also take a look at [their QuickStart guide](#), which does a good job of identifying the process to implement a metrics program. Yes, consensus involves lowest common denominators, and their metrics are no different. But keep things in context: to bootstrap any

metrics effort, organizations need to walk before they run, and in this case a lowest common denominator is exactly what's needed. Think of the consensus metrics as a common set of security metrics to start inter-company comparisons (benchmarks) now and to build upon over time. Taking a look at the CIS consensus metrics:

- **Incident Management:** Cost of incidents, Mean cost of incidents, Mean incident recovery cost, Mean time to incident discovery, Number of incidents, Mean time between security incidents, Mean time to incident recovery
- **Vulnerability Management:** Vulnerability scanning coverage, % systems with no severe vulnerabilities, Mean time to mitigate vulnerabilities, Number of known vulnerabilities, Mean cost to mitigate vulnerabilities
- **Patch Management:** Patch management coverage, Devices in compliance with patch policy, Mean time to patch, Mean cost to patch
- **Configuration Management:** % of configuration compliance, Configuration management coverage, Current anti-malware compliance
- **Change Management:** Mean time to complete changes, % of changes with security review, % of changes with security exceptions
- **Application security:** # of applications, % of critical applications, Application risk access coverage, Application security testing coverage
- **Financial:** IT security spending as % of IT budget, IT security budget allocation

Obviously there are many other types of information you can collect – particularly from your identity, firewall/IPS, and endpoint management consoles. Depending on your environment these other metrics may be important for operations. We just want to provide a rough sense of the kinds of metrics you can start with.

For those gluttons for punishment who really want to dig in we have built Securosis Quant models that document extremely granular process maps and the associated metrics for [Patch Management](#), [Network Security Operations](#) (monitoring/managing firewalls and IDS/IPS), and [Database Security](#).

We don't claim all these metrics are perfect. They aren't even supposed to be – nor are they relevant to all organizations. But they are a *place to start*. And the most important thing is to get started.

Qualitative 'Metrics'

I'm very respectful of Andy's work and his (correct) position regarding the need for any metrics to be numbers and have units of measure. Some things aren't metrics strictly speaking, but can still be worth tracking, and useful for benchmarking yourself against other companies. We call these "qualitative metrics," even though that's really an oxymoron. Keep in mind that the actual numbers you get for these qualitative assessments aren't terribly meaningful, but the trend lines are.

Having some context on your organization's awareness and attitudes around security is critical.

- **Awareness:** % of employees signing acceptable use policies, % of employees taking security training, % of trained employees passing a security test, % of incidents due to employee error
- **Attitude:** % of employees who know there is a security group, % of employees who believe they understand threats to private data, % of employees who believe security hinders their job activities

But any security program is about more than just the technical controls – a lot more. So qualitatively understanding the perception, knowledge, and awareness of security among employees is important.

We know what you are thinking. What a load of bunk! And for gauging effectiveness you aren't wrong. But any security program is about more than just the technical controls – a lot more. So qualitatively understanding the perception, knowledge, and awareness of security among employees is important. Not as important as incident metrics, so we suggest focusing on technical controls first. But ignore personnel and attitudes at your own risk, since the success of any security program also depends on perception and relevance in solving business-level problems.

Again, entire books have been written about security metrics. Our goal is to provide ideas and references for you to understand what you can count, but ultimately what you *do* count depends on your security program and business imperatives. Next we will focus on how to collect these metrics systematically. Because without your own data, you can't compare anything.

Collecting Data Systematically

Once you have figured out what you want to count (security metrics), the next question is how to collect the data. We look for metrics that are a) consistently and objectively measurable, and b) cheap to gather. That means some things we *want* to count may not be feasible. So let's go through each bucket of metrics and list the places we could get the data.

Quantitative Metrics

These metrics are pretty straightforward to collect, of course making the huge assumption you are already using some management tool to handle the function. That means the consoles that run operations like patching, vulnerabilities, configurations, and change management. Automation is really the linchpin of any metrics program, as having a common set of metrics gathered over time really sets the stage for operational improvement and inter-company comparisons to validate that improvement. So not only are there clear operational efficiencies to be gained through automating common functions, you also position your security program to quantify its performance.

- **Incident Management:** These metrics tend to be generated as part of the [post-mortem/quality assurance step](#) after closing an incident. Any post-mortem should be performed by a team, with the results communicated up the management stack; this means you should have consensus/buy-in on metrics such as incident cost, time to discover, and time to recover. We are looking for numbers with official units (like any good metric).
- **Vulnerability, Patch, Configuration, and Change Management:** These kinds of metrics should be stored by whatever tool you use for that specific function. The respective consoles should provide reports that can be exported (usually in XML or CSV format). Unless you use a metrics/benchmarking system that integrates with your tool, you'll need to map its output into a format you can normalize, and use for reporting and comparing to peers. But make sure each console gets a full view of the entire process, including remediation. Be sure every change, scan, and patch is logged in the system, so you can track the (mean) time to perform each function.
- **Application Security:** Metrics for application security tend to be a bit more subjective than we'd prefer (like % of critical applications developed securely), but ultimately things like security test coverage can be derived from whatever tools are used to implement the application security process. This is especially true for web application security scanning, QA, and other processes that tend to be tool-driven – as opposed to more amorphous functions such as threat modeling and code review.
- **Financial:** Your CFO and finance team is the place to start collecting financial data because they have metrics on what you spend. You can gather direct costs such as software and personnel, but indirect

costs are more challenging. Depending on the sophistication of your internal cost allocation, you may have very detailed information on how to allocate shared overhead, but more likely you will need to work with the finance team to estimate. Remember that precision is less important than consistency. As long as you estimate the allocations consistently, you can get valid trend data — but if you're comparing to peers you need to be a bit more careful about definitions.

For the other areas we mentioned, including identity, network security, and endpoint protection — this data will be stored in the respective management consoles. As a rule of thumb, the more mature the product (think endpoint protection and firewalls), the more comprehensive the data. And most vendors have already had requests to export data, or built-in more sophisticated management reporting/dashboards for large scale deployments.

But that's not always the case — some consoles make it harder than others to export data to different analysis tools. These management consoles — especially the Big IT management stacks — are all about aggregating information from many places, not necessarily integrating with other analysis tools. So as your metrics/benchmarking efforts mature, a key selection criterion will be the presence of an open interface to get data in and out.

Qualitative Metrics

Qualitative metrics are squishy by definition, and as such do not meet the requirements for 'good' metrics. The numbers on awareness metrics should reside somewhere, probably in HR, but may well not be aggregated. And percentage of incidents due to employee error is clearly subjective — assessed as part of the incident response process and stored for later collection. We recommend including that judgement as part of the general incident reporting process and captured like any other metric.

Attitude is much squishier — basically you ask your users what they think of your organization. Yes, you are asking about their perception of security. The best way to do that is an online survey tool. Tons of companies offer online services for that (we use [SurveyMonkey](#), but there are plenty). Odds are your marketing folks already have one you can piggyback on, but they aren't expensive. You'll want to survey your employees at least a couple times a year and track the trends. The good news is the companies all make it very easy to get the data out.

You can also think about how quantitative/qualitative hybrid analysis can provide the basis for interesting analysis. For example, correlating password resets or lost laptops within 3 weeks of security training can provide some insight into the effectiveness of the training, and thus the impact on overall security posture.

Systematic Collection

This is the point in the series where we remind you that gathering metrics and benchmarking are not one-time activities. It's part of an ongoing adventure. So you need to scope out the effort as a repeatable process, making sure you have the necessary resources and automation to collect this data *over time*. Collecting metrics on an *ad hoc* basis defeats the purpose, unless you are just looking for a binary (yes/no) answer. You need to collect data **consistently and systematically** to get real value.

Without getting overly specific about data repository designs and the like, you'll need a central place to store the information. That could be as simple as a spreadsheet or database, a more sophisticated business intelligence/analysis tool, or even an online service designed to collect metrics and present data. Obviously the more specific a tool is to security metrics, the less customization you'll need to generate the dashboards and reports needed to use these metrics as a management tool.

Sharing Data Safely

The best description of a security benchmarking effort we've seen is in Chapter 11 of [The Pragmatic CSO](#), which provides a good perspective on why benchmarking is important.

Since it is very hard to have objective, defensible measures of security effectiveness, impact, etc., a technique that can yield very interesting insight into the performance of your security program is to compare it to others. If you can get a sample set of standard questions, then you can get a feel for whether you are off the reservation on some activities and out ahead of others.

Benchmarking has been in use in other IT disciplines for decades. Whether on data center performance or network utilization, companies have always felt compelled to compare themselves to others. In fact, making a case to senior management for just about anything gets much better if you can say the magic words: "The other guys are doing it."

One of the best ways to figure out how good your security is, and get a feel for various other operational aspects of your security program, is to figure out how you compare to someone else. In fact, it's really the only truly objective way to evaluate effectiveness. The objective here is not to come up with a "security number" or "risk score", but to present information in the context of other companies that face the same kinds of attacks. This provides management with what they always want: a perspective on the level of risk they are willing to take, relative to the risk that other companies are taking.

The objective here is not to come up with a "security number" or "risk score", but to present information in the context of other companies that face the same kinds of attacks.

If you are behind a reasonable peer group, management may decide to invest more or to accept the risks of a less effective security program. If you are ahead, maybe they will opt to maintain or even accelerate investment ([in the unlikely event they can differentiate on security](#)). Or, yes, senior management might decide to scale back on security 'overhead'. Either way it's a win for you as the practitioner, because you know where you stand and the decision makers are actually making informed decisions based on solid data. How novel!

But before we can start thinking about comparing all the metrics we have decided are important and are now collecting systematically, we need some kind of infrastructure and mechanism to share this data safely and securely.

Throughout our research into building a security benchmark, it was clear that customers would require any sharing mechanism to ensure:

1. **Anonymity:** First and foremost, these customers wanted to make sure the data wasn't attributed back to them. No way, no how. Of all the things I discussed with these customers, this was non-negotiable. There could be no way for another customer to identify source data or derive which company provided any of it.
2. **Integrity:** The next issue was making sure the data was meaningful. That means it must be objectively and consistently gathered. Obviously there must be some level of agreement on what to count and how to count it, and that would likely fall under the purview of a third party.
3. **Security:** This goes hand in hand with anonymity, but it's different in that potential customers need to understand how the data would be protected (at a granular level) before they'd be comfortable sharing.

Given all that, is it any wonder that security benchmarking remains in its infancy? When talking to any potential community aggregator or commercial benchmark offering, be sure to dig very deeply into how the data is both secured and aggregated to calculate the benchmarks. You need to ensure proper data encryption and segregation to make sure your data doesn't get mixed with others, and that even if it somehow does, it wouldn't be accessible. Additionally, you'll want to make sure any device uploading data (this must be systematic and automated, remember) is mutually authenticated and authorized so no one can game the benchmark.

From an infrastructure protection standpoint, make sure *all* the proper controls are in place.

From an infrastructure protection standpoint, make sure *all* the proper controls are in place. Things like strong identity management, egress filtering, HIPS (if not whitelisting on all devices with access to the data), as well as significant monitoring on the network and database. Given some recent high-profile breaches, it's not unreasonable to expect network full packet capture as well. Ultimately you need to be comfortable with how your data is protected, so ask as many questions as you need to.

From an application standpoint it's also reasonable to expect the code to be built using some kind of secure development methodology. So learn about the threat models the vendor (or community) used to design the protection, as well as to what degree automated *and* non-automated testing mechanisms were used to scrutinize the application at all points during the development process. Learn about audits and pen tests, and basically crawl into very dark places in the provider's infrastructure to get comfortable.

This is a tall order and adds substantially to the due diligence required to get comfortable participating in a security benchmark. We understand this will be too high a hurdle for some. But keep your eyes on the prize: making security decisions based only actual data, within the context of your peer group. As opposed to doing what your gut tells you, or politics, or prayer.

Defining Peer Groups and Analyzing Data

Now we need to start deriving value from the data. Remember, metrics and numbers aren't worth the storage to keep them, if you don't use them as *management tools*. You need to start comparing the data, drawing conclusions, and adjusting your security program based on the data. In reality, objective measurements can only be put in context relative to another organization. OMG, actually making changes based on data rather than shiny objects, breaches, airline magazine articles, and compliance mandate changes? Radical!

Thus we first figure out what relative means, which involves defining peer group(s) for comparison. The first group you'll compare your data to is actually yourself. Yes, this is trend analysis on your own metrics. It provides some perspective on whether you are improving – but improving against yourself does not provide insight into whether you are spending too much money or focusing on the right stuff. This is where you need to think about benchmarking, or going beyond security metrics.

Peer Groups

There are several ways to define a peer group:

- **Industry:** This is your vertical market. Initially (until you have access to loads of data), you will focus on big industry buckets – defense, healthcare, financial, hospitality, etc. Obviously there are differences between investment banks and insurance companies within the financial verticals, but businesses in the same category have many consistent business processes which involve collecting very similar types of data. These organizations also tend to have similar geographic profiles – for example a typical retailer has a headquarters, regional distribution centers, and tons of stores. Additionally these companies exist under similar compliance/regulatory regimes. They also tend to be relatively consistent in terms of technology adoption and maturity, which is critical for making relevant comparisons.
- **Company size:** Similar to the consistencies we find among companies in the same vertical/industry, we also find many similarities between companies of roughly the same size. For instance large enterprises (10,000+ employees) are generally global by definition – it is very difficult to get that big while focusing on a single geographic region. So organizational models and scale tend to be fairly consistent within a company-size segment. These companies also tend to spend similarly on security. Of course there are always outliers and some industries show less consistency, but we aren't looking for perfection here.
- **Region:** Regional peers enable many interesting comparisons. Culture and attitudes toward security can be enhanced or hindered by government funding and compliance regimes. We also see relatively

consistent technology maturity/adoption within regions – largely based on local drivers such as compliance with laws and other rules, infrastructure, and available talent.

Of course, not all metrics apply to any peer group. So factor this in when you define your benchmark peer groups. Best to first figure out how the specific metrics correlate for each peer group. We know, it's math, but you'll figure out pretty quickly whether there are useful patterns or consistency within any particular metric. Focus on the metrics with the best correlation across a peer group.

Sample Size

Now that we're talking about math we need to address sample size. That's basically how much data you need before the benchmark is useful. And as usual *it depends*, but push for [statistical significance](#) over the long term. Why? Because by definition statistical significance means a result is unlikely to occur by chance. You don't want to make decisions at random, so that's our goal. More precisely, you want to *stop* making decisions based on chance.

But it's likely to take some time to get to a statistically significant dataset, so what can you do in the mean time? Look at the distribution, remove the outliers (which screw up your trend lines), and start comparing yourself against the trends you can spot. You can get a decent trend with only a handful of data points, for metrics that correlate strongly. Though remember this entire process needs to be based on candor and brutal honesty (even with yourself), so don't be tempted to massage the data to remove uncomfortable or confusing truths from the data. Always remember to keep the goal clearly in focus: to identify gaps and highlight success — neither of which requires a huge amount of data. But over time you are looking for statistical significance.

Reverting to the Mean

Another issue is whether you want to “revert to the mean,” meaning look like everyone else in your peer group. Once again, it depends. Let's look at a couple likely metrics categories:

- For spending, it's unlikely that you can get a reasonable return from security spending 3 standard deviations above the mean. Not unless your product/offering differentiates on security, [which is rare](#).
- For incidents, you want to push for perfection. Why? Because all your years of hard work can be unwound with one high profile breach. So the more effectively and quickly you respond and contain the damage, the better. The value in benchmarking your incident response comes from being better than the peer group, which can be used to win (or lose) points with senior management. You *definitely* don't want to be in the bottom quartile, which indicates a failure of incident response.
- For efficiency, effectiveness, and coverage metrics (most of the easily quantifiable and operational metrics), you want to be better than the mean. That shows operational competence.

In terms of prioritization, your spending is usually the most visible (to the folks who pay the bills, at least), so be in the ballpark there. Incidents come next, as they have a direct impact on issues such as availability and brand damage. Then comes the operational stuff – how you run the security program is critical to you and your team, but rarely interesting to senior management.

Communication Strategies

The stark reality is that many senior security folks are not comfortable talking to senior management and the Board of Directors. Thus, it tends not to happen frequently enough with the net result being a huge disconnect of both success and failure up and down the management stack to key security stakeholders. In fact, [the Pragmatic CSO](#) methodology originated largely to help technical folks figure out how to deal with their management responsibilities. The inability to communicate to key stakeholders will absolutely kill a benchmarking program because benchmarking entails ongoing incremental effort to gather metrics, as well as to compare against benchmarks and perform analysis. The benchmark must provide additional value, which must be communicated in order to make the effort worthwhile.

As we all know, nothing really happens by itself. You need to build a *systematic communications/outreach effort* to leverage the benchmark data, specifically targeting a number of constituencies important to the success of any security practitioner. Let's dig into how that's done, because it's a critical success factor for any benchmarking initiative.

Understanding Your Audience

The first rule of communications is to do it consistently and repetitively: tell them what you are going to say, say it, and then tell them what you just said. It sounds silly, but in today's over-saturated environment, where the typical C-level exec has a remarkably large number of things to worry about, you have no choice. Effective communication requires more than just talking a lot – you need to tailor your message to the audience. This is something security folks have always stunk at. If you have ever uttered the words “AV coverage” or “firewall rules” in a management meeting you know what I mean.

Senior management

If there is one thing to appreciate about senior management, it's that they are fairly predictable. Their interests involve things that directly impact revenues/expenses. Period. They don't want to know the details of how you do something unless it's off the rails. They want to know the bottom line and whether/how it will impact their ability to get paid their full bonus at the end of the year.

So for them we focus on incident data and budget efficiency. They want to know whether incidents have impacted availability and thus cost them money. They need to know about disclosures, with an eye toward brand damage. And they need to know how you do relative to peers, which provides the context for them to understand good and bad performance.

Getting time with senior folks is challenging. So you'll be doing well if you can get quarterly face time to go through the metrics/results/benchmarks. At a minimum you need to make your case annually before budgeting, but that is not really enough to get sufficient attention to successfully execute on your program.

Finally, how can benchmark data help you with these folks? Depending on your industry, senior management may be somewhat conservative, so having data about other companies doing it (whatever *it* may be), may eliminate whatever remain obstacles remain. It's an ugly job, but peer pressure does have value at times.

CIO

Odds are you report up through the technology stack, which means you'll spend some time with the CIO. This is a good thing, but keep in mind that the CIO's primary goal is to make sure technology is creating a competitive advantage (or at worst, not resulting in competitive disadvantage). We all know that security issues makes everyone look very bad. So focus on what interests senior management: incidents and budget efficiency. But with the CIO you should add high-level operational trending data, which highlights issues and/or shows progress on efficiency. Given the spend on security, the CIO needs to pay attention to efficiency and improve it.

How often should you be communicating with the CIO? Hopefully monthly, if not more often. We know it's hard to book time given the demands of your typical enterprise CIO. But you still need access and face time to make sure there is a clear understanding of where the security program is and what needs to be addressed.

Benchmark data helps substantiate the need for specific projects/investments, driven either by peer group adoption or efficiency/effectiveness gaps. Again, *your opinion* about what's important and needed is interesting, but not necessarily relevant. Data to substantiate your arguments makes them much stronger.

IT Ops teams

Likewise your pals in IT Ops don't want to be the reason for downtime or even worse, a data breach. You need their support to execute on any kind of security program, because Ops can help or hinder your protection efforts — with many different potential ways of creating problems for you and the CIO. But Ops isn't interested in the same metrics as senior managers, they worry about operational data, not financial. You need to focus discussions with them on areas where changes or activities depend on operational resources. As with all things operational, it's about increasing efficiency and reducing downtime due to error, so we want data which highlights issues, gaps, and areas to improve.

Ops folks may not appreciate being told they need to do things differently. This is another place where benchmark data can be your ace in the hole. By showing relative performance and ability to execute on operational processes, data can substantiate your arguments and help avoid having to go back to the CIO to complain "Ops sucks and makes our life hard!" — hoping the CIO will make them play nice.

Security team

Valuable as benchmark data is for telling a better story to stakeholders and key influencers of the security program, it is also a key management tool for your own security team. We all want our groups to work better and improve continuously — as we will discuss next. But without milestones and success criteria, and objective data to track progress, it's hard to demonstrate what needs to change or why, and to motivate your team to improve. Unless everyone has proper access to the data it's hard to keep them focused. We are big fans of open access to data, so we advocate making operational dashboards available to the team and/or scheduling periodic reports to communicate relevant data to the team on an ongoing basis.

Auditors

You thought we'd forget about your friend the auditor? Not so much. Keep in mind that auditors are different animals – they don't care about relative benchmarks. *They are entirely focused on their checklists.* So part of what you need to do is package up data to make it easier for the auditor to check his/her boxes. How? Your security metrics (suitably abstracted) substantiate the controls you have in place, and provide a gauge for their effectiveness. So whether you focus on those metrics or the comparative reports, it's about convincing the auditor you are in control of your program.

This isn't the only place benchmark data will be very handy. Sometimes you will have a "difference of opinion" with the auditor (we know – it's shocking). Showing that other companies are doing the same thing (or not) and/or showing that operationally you perform well compared to your peer group helps make your case.

Failure to Communicate Is Not an Option

We can't stress enough the importance of having a structured communications plan to discuss the metrics/benchmark data with all the appropriate stakeholders early and often. This doesn't happen by itself, so you need to be persistent – senior folks have been known to cancel meetings with the security team – and diligent about getting your time and making the most of it.

Continuous Improvement

Now it's time to shift focus internally. One of the cool things about security metrics and benchmarks is the ability to analyze trends over time and use that data to track progress against your key goals. Imagine that – managing people and programs based on **data**, not just gut feel.

Besides being able to communicate how you are doing on security much more authoritatively, you can also focus on *continuously improving* your activities. This is a good thing to do – particularly if you want to keep your job. We will talk about the importance of consistency in gathering data and benchmarks over a long period of time, and then how to get ongoing value from benchmarks by using them to mark progress toward a better and more secure environment.

Besides being able to communicate much more authoritatively how you are doing on security, you can also focus on *continuously improving* your activities.

Programs and Feedback Loops

We don't want to put the cart ahead of the horse so let's start at a high level, with how to structure the security program so it's focused on improvement rather than mere survival. Here are the key steps:

- Define success (and get buy-in up the management stack)
- Determine success characteristics and break down into activities that will result in success
- Quantify those activities, determine appropriate metrics (focus on metrics impactful to business operations), and set goals for those metrics (define success)
- Set objectives for each activity and communicate those objectives
- Run your business; gather your metrics
- Analyze metrics; report against success criteria/objectives (can be reported as a % of goal attainment to keep the goal in constant focus)
- Identify gaps, address issues, and reset objectives accordingly
- Lather, rinse, repeat (weekly, monthly, etc.)

Digging deeply into security program design and operation would be out of scope, so we'll just refer you to Mike's methodology for building a security program: [The Pragmatic CSO](#).

Communicating to the Troops

We talked about communicating with key stakeholders in the security process, and your security team is clearly a primary constituency. So let's revisit that discussion and its importance. Your team needs to understand the process, how benchmark data will be used to determine success, and what the expectations will be.

Don't be surprised to experience some push-back on this new world order, which might be quite significant. Put yourself in your team's shoes for a moment. For most of these folks' careers they have been evaluated with squishy subjective assessments of effectiveness and effort. Now you want to move them to something more quantified, where *they can neither run nor hide*. Top performers should not be worried – at all. That's a key point to get across.

Exercise some patience in getting folks' heads in the right spot, but remember that you aren't negotiating. Part of the justification for investing (rather significantly) in metrics and benchmarks is to leverage that data in operations. You can't do that if the data isn't used to evaluate performance – both good and bad.

It's Not a Tool, It's a Lifestyle

Another point to keep in mind is that this initiative isn't a one-time thing. It's not something you do for an assessment, and then forget in a drawer the moment the auditor leaves the building. Benchmarking, done well, becomes a key facet of managing your security program. This data becomes your North Star, providing a way to map out objectives and stay on course to reach them. We have seen organizations start with metrics as a means to an end, and later recognize that they can change everything about how operational efforts are managed, perceived, and supported within the organization. The lack of security data has hindered acceptance of benchmarking in the security field, but it's time to revisit that.

As usual, there are some caveats to data-driven management. No one size fits all. We see plenty of cultural variation, which may require you to take a less direct path to the benchmark promised land. But there can be no question about the importance of quantifying activity, effectiveness and efficiency compared to *not* quantifying it.

Looking Inward

We have spent much of this paper on why benchmarking is important. But we also need to point out some situations where benchmarking may *not* be appropriate. There are clearly situations where you can't benchmark, particularly on granular operational data, which I call *Ninja Metrics*.

Dependency: Peer Group Data

Most organizations have 'nascent' metrics programs, which may actually be too kind. But not all. Some have embraced detailed programs that gather all sorts of data, mostly focused on operations. This represents the next step of a metrics program, and our suggestions for how to proceed are available in our various *Quant* research projects. We have created highly granular process maps (with associated metrics) for [Patch Management](#), [Network Security Operations](#), and [Database Security](#). Each report specifies 50+ distinct metrics you can measure for that discipline. Yes, they are comprehensive.

The key dependency in implementing a benchmarking effort is the availability of peer data for comparison.

But there is an obvious issue with benchmarking at this level. You will have a hard time finding similarly granular data from other companies for comparison. So the key dependency in implementing a benchmarking effort is the availability of peer data for comparison. We hope that over time enough companies will start tracking granular operational metrics, and become comfortable enough with benchmarking to share their data, perhaps aided by the emergence of commercial benchmarking offerings. In the meantime you can (and should) continue to push your

metrics program forward, which may require bootstrapping an effort within the construct of an industry group.

Compare against Yourself

What do world class athletes do when they reach the top of the heap? You know, folks like Michael Phelps, who has basically shattered every record there is to shatter. They start comparing themselves to their past performance. Improvement is measured internally rather than externally. Even if no one else has ever done better, you know you can. And this is what you will likely need to do the most granular operational functions.

When you take a step back this makes a lot of sense. The reality is that you aren't necessarily trying to 'win' on some arbitrary standard of operational excellence. You want to *improve*. That said, it is important to have an idea of where you stand in comparison to everybody else, at least on the high-level operational metrics. But for the most granular metrics, not so much.

Benchmarking in Action

Let's wrap up by actually walking through a scenario and applying the process. Yes, it's a bit contrived, but we'll hit the high points of the process, deciding where to start, collecting the data, establishing the peer group, and communicating findings. Keep in mind that we focus on *quick wins* to show immediate value, build momentum, and leverage that momentum for programmatic success.

Scenario

Let's use a mid-tier financial company as our example. I'd pick a large enterprise, but there are many nuances and moving pieces within a large enterprise that would require more detailed discussion. So let's keep it relatively simple. Likewise, we picked the financial vertical because of 1) need and 2) availability of data. Financial regulatory oversight has created a general perspective of security first, and data-centricity (yes, these are the folks who try to do risk management for a living) means these businesses are more likely to embrace a benchmarking mentality.

In our (contrived) scenario, the board drove hiring a new CISO to "fix security." As easy as it is to see this as just catering to a board directive, the senior team seems to have a commitment to fix things and do it the right way. So the CISO has a clear honeymoon period and some leeway in thinking somewhat unconventionally about how to build the security program. The new CISO still spends some time figuring out what's installed and what's not working, but he knows the organization has AV deployed, they use an external scanning service, and they do a pretty good job of patching internal systems.

But like many smaller financial institutions they use hosted applications for most of their business processing. So a lot of their data is outside their direct control. Over the past few years the organization has had a handful of incidents, but none really resulted in major data loss. So the CISO was pleasantly surprised when he got a mandate to fix the security program, despite it not being obviously broken. The senior team concluded they were living on borrowed time and wanted to act decisively to ensure they were ready when the brown stuff hit the fan (which it inevitably does).

See? We told you the scenario was contrived, but without a senior-level mandate to make changes and implement a security program, getting any kind of security metrics/benchmarking initiative going is quite difficult (if not impossible).

Where Do You Start?

Now the CISO has to figure out where to start. He has decided to find his most visible gaps first. You know, the ones you could drive a Mack truck through. So he starts with a comprehensive risk assessment to build a baseline, but he also wants to compare his environment to other like-sized companies (both inside and outside his industry) to figure out how he compares to those organizations. Boiling the ocean and trying to

do everything at this point would be a bad idea. He'd get buried in the nuances of the data and not get anything done, which could endanger his entire security program. So he needs to ask the following questions:

- **What do you need to achieve? Where are the key operational problems?** This is where you always need to start. In our case study, the CISO is looking to identify his most critical gaps, and given the *luck* they've had in not having a huge data loss even with a few breaches, he wants to start with incident response.
- **What data do you have?** Next you have to figure out if you have the data or can get it easily. With incident data, the reality is typically that findings from forensics investigations are available, but haven't been collected or formatted for comparison. *But the data exists*, so it makes sense to keep pressing down this path. If the data doesn't exist or can't be gathered quickly, it's time to look at Plan B. You don't want to hold up the effort — right now you need a quick win.
- **Which comparative data will be most convincing?** The initial focus on incident response represents a pretty shrewd move for the new CISO. He knows the board and senior management are sensitive to getting nailed, and a set of reasonable consensus metrics are available (from [CIS](#)), and he has the data. This all increases his chance of success.

Peer Groups and Service Providers

Next, our CISO has to define the peer group for analysis. This isn't brain surgery. He'll need to compare to other financials (duh!), but also companies in other regulated industries (such as healthcare and utilities) of similar size. The good news is there are a ton of mid-sized hospital groups, as well as many community utilities, with similarly sensitive data. But how do they get their hands on it for comparison?

Now we go back and revisit the selection criteria for any kind of provider you'd think about for benchmarking services. Remember, these folks have to 1) have access to the data you'd need and 2) be able to protect the data you share with them. You may not be able to get everything done with any one provider. In our case study here, the CISO actually picks two. The first is his regional bank ISAC, which has been gathering data from members for a while. The second is a commercial benchmarking offering, which has more data about other industries that aren't the focus of the ISAC. The CISO would prefer a single provider, but until a critical mass of data for many verticals is captured, he'll need to piecemeal the solution to solve the problem.

Analyze

Equipped with data regarding his first area of focus (incident response), the CISO can now start to analyze the results to find the biggest gaps in his process. The data shows that the board's instinct that they've been lucky is right on the money. Most organizations in the finance peer group have had more breaches and significant loss from at least one. Additionally, comparably sized companies have also had more breaches, but on the good news front, those comparable incidents took those companies longer to handle. So what conclusions does the CISO draw? That their program isn't terrible, but they need to do more to ensure any incidents can be handled faster and more effectively.

Luck is not a strategy, so there are clear areas for improvement. This will require looking at some enhanced monitoring technologies to detect potential issues faster and a network full packet capture capability to identify root causes and enable more precise forensic analysis. He also wants to invest in some forensics training for his team, and get a top-notch incident response firm on a small retainer to make sure he knows who to call when something goes down. Equipped with this data, he can make a more compelling case for the new equipment and services. A quick win in the bag.

What's Next?

After finishing up with one area, the CISO can focus on the next area for metrics/benchmark analysis. Given the consistent data from industry breach reports, he decides to focus on patching effectiveness. He understands that many of the breaches resulting in huge amounts of lost data started with attacks where a patch already existed, but wasn't applied. Again, the CISO has a patching program in place, but without an idea about its relative performance, he's shooting in the dark. He could have just as easily decided to think about vulnerability management or firewall operations. It doesn't matter what he picks, so long as the data aligns with areas needing operational improvement and representing areas of potential exposure.

Summary

Given the lack of objective evidence regarding the performance of any security program, we believe the idea of security benchmarking is critical to substantiating the impact of the millions of dollars the typical company spends on security each year. As mentioned numerous times during this paper, benchmarking and the underlying metrics programs are not quick fixes, rather requiring a long term commitment to collecting data and doing the requisite analysis.

But if you've ever sat in a meeting with senior management and not been able to convince them your efforts are worthwhile. Or if you've ever gotten your budget slashed because you couldn't make the case as to why you needed to implement a new control, then you've felt our pain. We believe the best way to start addressing the issue is to quantify your efforts and compare those metrics with a clear and defined peer group. All of us at Securosis wish you luck on your journey.

If you have any questions on this subject matter, 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 a 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, 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 assessments.
- *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. Example services available as part of a retainer package include market and product analysis and strategy, technology guidance, product evaluations, and merger and acquisition assessments. 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 services tend to be customized to meet a client's specific 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/>.