Dynamic Security Assessment

Version 1.5
Released: March 27, 2017

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

Author’s Note

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

Special thanks to Chris Pepper for editing and content support.

This report is licensed by SafeBreach.

www.safebreach.com

SafeBreach helps answer the question security leaders are being asked today. “Are we secure?” and “Can our defenses stand up to a breach?”

Our platform simulates hacker breach methods across the entire kill chain to quantify risks and validate security controls, using an extensive and growing Hacker’s Playbook™ of research and real-world investigative data.

SafeBreach delivers a simulation platform for the “Dynamic Security Assessment” process defined in this paper. For more information, visit www.safebreach.com or follow on Twitter @SafeBreach.

Copyright

Dynamic Security Assessment

Table of Contents

The Limitations of Security Testing

We have been fans of testing the security of infrastructure and applications, at least since we started doing research. We have always known attackers are testing your environment all the time, so if you aren't also self-assessing, inevitably you will be surprised by a successful attack. Like most security folks, we're no fans of surprises.

Security testing and assessment has gone through a number of iterations. It started with simple vulnerability scanning. You could scan a device to understand its security posture, which patches were installed, and what vulnerabilities remained. Vulnerability scanning remains a function at most organizations, driven mostly by a compliance requirement.

As useful as it is to understand which devices and applications are vulnerable, a simple scan provides limited information. A vulnerability scanner cannot recognize that a vulnerable device is not exploitable due to other controls. The discipline of penetration testing emerged to go beyond simple context-less vulnerability scanning, with trusted humans trying to steal data.

Pen tests are useful because they provide a sense of what is *really* at risk. But a penetration test is resource-intensive and expensive, especially if you use an external testing firm. To address that we used automated pen testing tools, which use actual exploits in a semi-automated fashion to simulate an attacker.

With the easy availability of penetration testing tools (notably the open source Metasploit), defending against a pen testing tool has emerged as the low bar of security. Basically, if you cannot stop a primitive attacker using Metasploit (or another pen testing tool), you aren't very good at security.

Regardless of whether you use carbon-based (human) or silicon-based (computer) penetration testing, the results characterize your environment at a single point in time. As soon as you blink, your environment will have changed, and the validity of your findings starts to degrade.

With the easy availability of penetration testing tools (notably the open source Metasploit), defending against a pen testing tool has emerged as the low bar of security. Our friend Josh Corman coined HDMoore's Law, after the leader of the Metasploit project. Basically, if you cannot stop a primitive attacker using Metasploit (or another pen testing tool), you aren't very good at security.

Raising the Low Bar

As we lead enterprises through developing security programs, we typically start with adversary analysis. It is important to understand what kinds of attackers will target your organization and what they will be looking for. If you think your main threat is a 400-pound hacker in his parents' basement, defending against an open source pen testing tool is probably sufficient.

> The key thing to understand about adversaries is: They don't play by your rules.

But do any of you honestly believe an unsophisticated attacker wielding a free penetration testing tool is *all* you have to worry about? Of course not. The key thing to understand about adversaries is: *They don't play by your rules*. They will attack when you don't expect it. They will take advantage of new attacks and exploits to evade detection. They will use tactics that look like a different adversary to raise a false flag.

An adversary will do whatever it takes to achieve their mission. They can usually be patient, and will wait for you to make a mistake. So the low bar of security represented by a penetration testing tool is not good enough.

Dynamic IT

The increasing sophistication of adversaries is not the only challenge in assessing your environment and understanding risk. Technology infrastructure seems to be in the middle of the most significant set of changes we have ever seen, which dramatically complicates your ability to assess your environment.

First, you have no idea where your data actually resides. Between SaaS applications, cloud storage services, and integrated business partner networks, the boundaries of traditional technology infrastructure have been extended unrecognizably, and you cannot assume your information is on a network you control. And it is hard (or even illegal) to test a network you don't control.

The next major change under way is mobility. Between an increasingly disconnected workforce and an explosion of smart devices accessing critical information, you can no longer assume your employees will access applications and data from your networks. Authorized users need access to data from anywhere in the world, at any time, which seriously complicates assessment.

Finally, the push to public cloud-based infrastructure obscures where your compute and storage are. Many enterprises we work with are building cloud-native technology stacks using dozens of services across cloud providers. And as a bonus you don't necessarily know where you will be attacked either.

To recap, you no longer know where your data is, where it will be accessed from, or where computation will happen. And it is your responsibility to protect information in this dynamic IT environment, so you need to assess the security of your environment as frequently as practical. This is the challenge of security assessment today, with a hint at how much more complicated it will be tomorrow.

We Need Dynamic Security Assessment

The only way to keep pace with your dynamic IT environment is *dynamic* security assessment. The rest of this paper will lay out what means and how to implement it.

As an introduction, a dynamic security assessment offering includes:

- A highly sophisticated simulation engine, which can imitate typical attack patterns from sophisticated adversaries without putting production infrastructure in danger.
- An understanding of the local network topology, for modeling lateral movement and isolating targeted information and assets.
- Access to a security research team to leverage both proprietary and public threat intelligence, and to model the latest and greatest attacks to avoid unpleasant surprises.
- An effective security analytics function to figure out not just what is exploitable, but also how different workarounds and fixes will impact infrastructure security.

Process and Functions

Let's start fleshing out the concepts underlying Dynamic Security Assessment by laying out the requirements for our vision:

1. **Ongoing:** Infrastructure is dynamic, which makes traditional vulnerability testing less useful — a point-in-time assessment can be obsolete before the report hits your inbox.

2. **Current:** Every organization faces fast-moving and innovative adversaries, leveraging ever-changing attack tactics and techniques. To provide relevant and actionable findings a testing environment must be up-to-date, and the test must factor in new tactics.

3. **Non-disruptive:** The old security testing adage of "do no harm" still holds. Assessment functions must not take down systems or hamper operations in any way.

4. **Automated:** No security organization (that we know of, at least) has enough people, so expecting them to constantly assess the environment isn't realistic. To make sustained assessment feasible it must be mostly automated.

5. **Evaluate Alternatives:** Once a potential attack is identified you need to validate and then remediate it. Don't waste time shooting into the dark. You should be able to predict the impact of potential changes and workarounds to first figure out whether they would stop the attack in question, and then select the best option if you have several.

Dynamic Security Assessment Process

As usual we start by focusing on process rather than shiny widgets. Fortunately it's pretty straightforward.

1. **Deployment:** Your first step is to deploy assessment devices, which you might call agents or sensors. But you need presence both inside and outside the network to launch attacks and track results.

2. **Define Mission:** After deployment you need to figure out what a typical attacker would want to access in your environment. This might be a formal threat modeling process, or you could start with the simple question, "What could be compromised that would cost the CEO/CFO/CIO/CISO his/her job?" Everything is important to the person responsible for it, but to find an adversary's most likely or most important target, consider what would drastically harm your business.

3. **Baseline/Triage:** Next you need an initial sense of the vulnerability and exploitability of your environment, and a library of attacks to investigate its vulnerability. If you try you can probably identify the main critical issues which would immediately require all hands on deck. Once you get through initial triage and remediation of potential attacks, you will have an initial activity baseline.

4. **Ongoing Assessment:** Then you can start assessing your environment on an ongoing basis. An automated feed of new attack tactics and targets is useful for ensuring you look for the latest attacks seen in the wild. When the simulation engine finds something, administrators are alerted to successful attack paths and/or patterns for validation, and then determination of the criticality of a potential attack. This process needs to run continuously because things change from minute to minute.

5. **Fix:** This step tends to be performed by Operations, and is somewhat opaque to the assessment process. But it is where critical issues are fixed and/or remediated.

6. **Verify Fixes:** The final step is to validate that issues were actually fixed. The job is not complete until you verify that the fix is both operational and effective.

Yes, that all looks a lot like any other security assessment methodology. What needs to happen hasn't really changed — you still need to figure out exposure, understand criticality, fix, and then make sure the fixes worked. What changes is the technology used for assessment. This is where the industry has made significant strides to improve both accuracy and usefulness.

> What needs to happen hasn't really changed — you still need to figure out exposure, understand criticality, fix, and then make sure the fixes worked. What changes is the technology used for assessment.

Simulation Engine

The centerpiece of DSA is an *simulation engine*. It enables you to understand what is possible in an environment, to define the universe of possible attacks, and then to figure out which would be most damaging. This effectively reduces the detection window, because without it you don't know if an attack has been used on you. It also helps you prioritize remediation efforts, focusing on what would work against your defenses.

Analysis needs to be considered from two perspectives. The first we call *white box* testing: analysis of exposures in the environment, assessed with knowledge of the environment. This involves running a scanner, fuzzer, or some similar tool to check known exposures.

To start a white box test, you feed the topology of your network into the engine, because attackers need to first gain a foothold and then move laterally to achieve their mission. Once your engine has a map of your network, existing security controls are factored in so the engine can determine which devices are vulnerable to which attacks. For instance you'll want to specify access control points (firewalls) and threat detection (intrusion prevention) points in the network, and what kinds of controls run on which endpoints. Attacks almost always involve both networks and endpoints, so your engine must be able to simulate both.

The second assessment is a form of *black box* testing, where analysis is performed without any knowledge of the environment. You set up sensors on sensitive segments you *really* don't want attackers to access, as well as outside the network, and let the tool run through its battery of attacks to see what it can do. This is more like an automated penetration test, but the tool runs automatically and continuously to assess the environment more like an adversary would.

The magic is in how the simulation engine figures out *what* can be attacked and *how*. The best practices of attackers are distilled into algorithms to simulate how an attack could strike across multiple networks and devices. Consider the attack lifecycle/kill chain. The engine simulates hacker activity from both inside and outside your network to determine what is visible and where to move next in search of its target. Sometimes they get lucky and are invited in by unsuspecting employees, but other times they look for weaknesses in perimeter defenses and applications. Everything is fair game for data collection and DSA.

With an idea of which controls are active on each device you can determine which attacks might work. Using data from reconnaissance, an attack path from entry point to target can be generated. These paths represent lateral movement within your environment, and the magic of dynamic assessment lies in figuring out how an attacker would move without actual repercussions to your network.

Finally you will want to assess the ability of an attacker to exfiltrate data, so the assessment system will try to get the payload past egress filters.

But this isn't an either/or proposition. The correct answer is *all of the above*. DSA algorithms provide a probabilistic analysis of your attack surface, with white box texting to provide extensive coverage of everything within the infrastructure, and the black box to approximate what an adversary can and would do when gaining presence in your environment.

You cannot run constant penetration tests on everything, so Dynamic Security Assessment helps identify areas of concern; then you can have a human check and determine the most appropriate workaround.

It is not possible to fully mimic a human attacker presented with specific and changing defenses programmatically. That's what red teams and penetration testers are for. But you cannot run constant penetration tests on everything, so Dynamic Security Assessment helps identify areas of concern; then you can have a human check and determine the most appropriate workaround.

This multi-faceted approach helps you understand likely paths for attackers to access your targets and exfiltrate data. As we mentioned above, in software testing terms DSA increases coverage. Humans cannot consider every attack, try every path, and attack every device — so a DSA system can fill in your overview.

Threat Intelligence

Referring back to our requirements, the simulation/analytics engine handles most of what you need done. It provides ongoing, non-disruptive, automated assessment of your entire environment. The only thing missing is keeping the tool current, which is where threat intelligence (TI) comes into play.

Integration of new attacks into an simulation engine enables it to model new tactics and targets. If you face a sophisticated adversary you can get some idea of what they are likely to throw at you from other organizations' reports. You can feed the engine new methods to analyze. If a new attack would succeed, you should know about it *before* it succeeds in your environment.

Automation is essential for sustainable and useful assessment. You don't have time to manually keep the product updated with the latest attacks and run each new test. Many security operations folks resist automating functions because getting it wrong can mean downtime. But you have more leeway with assessment where a faulty update won't disrupt your environment. You might get some annoying false positives but you won't lose half your network as you could if a change to an active endpoint or network security control goes awry.

Visualization

Finally, once you have an attack that *could* succeed, you'll want to dig into specifics. The modern way to do that is visualization. You should be able to see an attacker's path and which devices could be compromised. Drilling down into specific devices, with possible attacks highlighted by the simulation engine, can help you identify faulty controls and weak configurations.

Visualization is key to weighing alternative fixes and figuring out which would be most efficient. Assessing how different controls would affect a simulated attack can help you quickly identify your best remediation option.

If dynamic security assessment sounds like what vulnerability management should have evolved into, you are right. a DSA engine puts vulnerabilities into context. It's not just about what can be attacked, but how each attack would fit into a larger campaign to access a target and steal information.

If dynamic security assessment sounds like what vulnerability management should have evolved into, you are right. Rather than looking at devices individually and providing summary data with dashboards showing how quickly you are fixing vulnerabilities, a DSA engine puts vulnerabilities into context. It's not just about what can be attacked, but how each attack would fit into a larger campaign to access a target and steal information.

DSA in Action

To illuminate these concepts and make things a bit more tangible let's consider a scenario involving a large financial services enterprise with hundreds of locations. Our organization has a global headquarters on the West Coast of the USA, and 4 regional headquarters across the globe. Each region has a data center and IT Operations folks to run things. The security team is centralized under a global CISO, but each region has a team to work with local business leaders to ensure proper protection and handle jurisdiction. The organization's business plan includes rapid expansion of its retail footprint and additional regional acquisitions, so the network and systems will continue to become more distributed and complicated.

New technology initiatives are being built in the public cloud. This was controversial at first but there isn't much resistance any more. Migration of existing systems remains a challenge, but cost and efficiency have steered strategy toward consolidation of regional data centers into a single location to support legacy applications within 5 years, with all other systems running in the cloud. This centralization is being enabled by moving a number of back-office systems to SaaS. Fortunately the back-office software provider just launched a new cloud-based service, which makes scaling the system for new locations and integration of acquired organizations much easier. Our organization has been using cloud storage heavily — since initial fears were overwhelmed by cost savings from reduced investment in the complex and expensive on-premise storage architecture.

Security is an area of focus and a major concern, given the quantity and sensitivity of financial data our organization manages. They are constantly phished and spoofed, and applications are under daily attack. There are incidents, fortunately none severe enough to require customer disclosure, but the fear of missing adversary activity is always there.

For security operations, they currently scan their devices and have a reasonably effective patching/hygiene processes, but still average 30 days to roll an update out across the enterprise. They also undertake an annual penetration test. To keep key security analysts engaged they allow them to spend a few hours per week hunting active adversaries and other malicious activity.

CISO Concerns

The CISO has a number of concerns regarding this organization's security posture. Compliance mandates require vulnerability scans, which enumerate theoretically vulnerable devices. But working through the list and making changes takes a month. They always get great information from the annual penetration test, but that only happens once a year.

Moreover, the scans and pen tests apply to just existing systems across current data centers. The move to the cloud is significant and accelerating. As a result sensitive protected data is all over the place, and they need to understand which ingress and egress points present what risk of both penetration and exfiltration.

Compounding the concern is the directive to continue opening new branches and acquiring regional organizations. The initial diligence on each newly acquired environment takes time the team doesn't really have, and they need to make security compromises to hit their aggressive timelines — to integrate new organizations and drive cost economies.

The bottom line is that their exposure window lasts at least a month, when everything works well. They know it's too long, and need to understand what they should focus on — accepting they cannot get everything done — and how they should most effectively allocate personnel.

In an attempt to get ahead of attackers, they undertake some hunting. But it's a part-time endeavor for staff, who tend to find easy stuff because that's what their tools identify first.

The bottom line is that their exposure window lasts at least a month, when everything works well. They know it's too long, and need to understand what they should focus on — accepting they cannot get everything done — and how they should most effectively allocate personnel.

Using Dynamic Security Assessment

The CISO understands the importance of assessment — as demonstrated by their existing scanning, patching, and penetration testing practices — and is interested in evolving toward a more dynamic assessment methodology. For them DSA might look something like the following:

- **Baseline Environment:** The first step is to gather network topology and device configuration information, and build a map of the current network. This data can be used to baseline traffic flow through the environment, along with what attack paths could be exploited to access sensitive data.

- **Simulation/Analytics:** This financial institution cannot afford downtime to their 24/7 business, so a non-disruptive and non-damaging means of testing infrastructure is required. They must be able to assess the impact of adding new locations and (more importantly) acquired companies to their own networks, and understand what must be addressed before integrating each new network. Finally, a cloud network presence offers an essential mechanism for understanding the organization's security posture, because an increasing amount of sensitive data has been and continues to be moved to the cloud.

- **Threat Intelligence:** Fortunately our model company is big, but not a Fortune 10 bank. So it will be heavily targeted, but not at the bleeding edge of new large-scale attacks using highly sophisticated malware. This provides a narrow window to learn from other financials — seeing how they are targeted; the malware used; the bot networks it connects to; and other tactics, techniques, and procedures (TTPs). This enables them to both preemptively put workarounds in place and understand the impact of workarounds and fixes under consideration before actually committing time and resources to specific changes. In a resource-constrained environment (any modern environment) this is essential.

Dynamic Security Assessment's new capabilities can provide a clear advantage over traditional scanning and penetration testing. The idea isn't to supplant existing methods, but to supplement them to provide more reliable means of prioritizing effort and detecting attacks.

Bringing It All Together

For our sample company the first step is to deploy sensors across the environment, at each physical location and within all the cloud networks. This provides data to model the environment and build the initial network map. With the environment modeled you can start analyzing risks to sensitive data stores. Identifying a handful of 'missions' adversaries are likely to undertake helps focus efforts on clear and present dangers, and avoid getting distracted or falling into every potential hole.

This initial assessment and resulting triage help focus efforts on attacks which can cause real damage. The CISO understands the 30-day window before things can be addressed, but the team can focus on issues with high-profile networks and devices which put sensitive data at risk.

Once initial triage is done, the team can undertake a more detailed analysis of the environment, turning the map into a baseline, understanding typical traffic flows and activities within all the organization's systems and networks. This improves both simulations and ongoing assessments to identify anomalous activity which warrants further investigation and/or immediate action.

Threat intelligence data also feeds into ongoing assessment enabled by DSA. Instead of just patching everything first-come-first-served, the CISO can marshal resources to address new attacks seen in the wild which would be work in this environment, as indicated by the ongoing simulation. Again, this helps focus resources on the most threatening issues.

DSA also helps with change control. As new changes are requested, driven by business needs and application upgrades, their impact can be modeled so risks are understood. When an application is deployed in the cloud, for instance, the network map can be updated quickly to reflect new potential exposure. Similarly, diligence on new offices and integration of acquired companies can be accelerated because new locations can be easily modeled and their risks evaluated. What used to be an unscientific *ad hoc* process can be quick and fact-based. This enables the CISO to present concerns along with hard data about potential risks, not just gut feel.

Finally, DSA ensures that changes are made completely and accurately. Ongoing assessment identifies issues which have been successfully vs. unsuccessfully remediated, and what else needs to be done if anything. The CISO is able to address the most serious concerns, to focus on the largest risks, and to get a full picture of the entire infrastructure — including resources now in the cloud.

Summary

The most critical success factor for almost every security program is its ability to prioritize activity and allocate scarce resources effectively. An invaluable technique for determining the areas of most risk is to assess your environment on an ongoing basis, understanding that sensitive corporate data is all over the place and new attacks appear daily. Being aware of all these moving pieces and simulating how existing defenses will work against these attacks gives you a view into the future, enabling you to tune your defenses *before* adversaries show you your weaknesses.

OK, it's not an actual view into the future, but it's as close as you can get in such a dynamic environment. We cannot underscore enough the importance of understanding the impact of changes to your environment *before* making them, given the severe resource constraints every organization faces today. If an assessment system can tell you which controls are likely to work and which aren't, you can reduce the time wasted patching systems which aren't exposed and deploying controls which won't improve your security posture meaningfully.

We cannot underscore enough the importance of understanding the impact of changes to your environment before making them, given the severe resource constraints every organization faces today.

As you look for leverage points to increase the efficiency and effectiveness of your security program we recommend considering Dynamic Security Assessment as a means to learn and experiment before you make changes. For most of us who grew up in the School of Hard Knocks, this is a very welcome change.

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

About the Analyst

Mike Rothman, Analyst and President

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

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

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

About Securosis

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

- **Primary research publishing**: We currently release the vast majority of our research for free through our blog, and archive it in our Research Library. Most of these research documents can be sponsored for distribution on an annual basis. All published materials and presentations meet our strict objectivity requirements and conform to our Totally Transparent Research policy.

- **Research products and strategic advisory services for end users**: Securosis will be introducing a line of research products and inquiry-based subscription services designed to assist end user organizations in accelerating project and program success. Additional advisory projects are also available, including product selection assistance, technology and architecture strategy, education, security management evaluations, and risk assessment.

- **Retainer services for vendors**: Although we will accept briefings from anyone, some vendors opt for a tighter, ongoing relationship. We offer a number of flexible retainer packages. Services available as part of a retainer package include market and product analysis and strategy, technology guidance, product evaluation, and merger and acquisition assessment. Even with paid clients, we maintain our strict objectivity and confidentiality requirements. More information on our retainer services (PDF) is available.

- **External speaking and editorial**: Securosis analysts frequently speak at industry events, give online presentations, and write and speak for a variety of publications and media.

- **Other expert services**: Securosis analysts are available for other services as well, including Strategic Advisory Days, Strategy Consulting engagements, and Investor Services. These tend to be customized to meet a client's particular requirements.

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

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