



Interested in learning
more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Need for Speed: Streamlining Response and Reaction to Attacks

Copyright SANS Institute
Author Retains Full Rights



Sponsored by McAfee

Need for Speed: Streamlining Response and Reaction to Attacks

May 2013

A SANS Whitepaper

Written by: J. Michael Butler

Why Speed and Accuracy Are Crucial PAGE 2

Applying “Real-Time” to Analytics PAGE 4

The Importance of Automated Response PAGE 6

What Should Security Analytics Do? PAGE 8

Better Analysis Leads to Better Security PAGE 9

Case Study in Efficiency Improvements PAGE 13

Special Considerations for the Cloud PAGE 14

Introduction

“Despite their investments in endpoint security systems, organizations are waking up to the ugly truth that they are nearly blind when it comes to advanced attacks and [malware](#) lurking in their networks.”¹ Advanced threats are getting into protected networks, spreading themselves to critical servers and avoiding detection by many of today’s security sensors. Because threats to endpoints have become increasingly advanced and are often sponsored and well-funded by malevolent organizations, we can no longer rely on traditional tools to protect ourselves. The zero-day threat has become more dangerous, enabling malicious actors to take advantage of vulnerabilities we may not yet know exist.²

Most often, threats enter the enterprise through compromised endpoints, as seen in the 2013 Data Breach Investigations Report, published by Verizon.³ Modern endpoints are increasingly complex in nature. Today’s average American business has at least one computer on every desk. Companies issue laptops to most of their knowledge workers, and many—if not most—employees have smartphones they use for work applications as well as risky applications of their own. Businesses are embracing social networking and the cloud, with the Internet tying businesses together in a way that was science fiction 40 years ago.

The ever-expanding number of “things” in the corporate IT environment has created new problems for the security practitioner. They connect everywhere, to each other and to the Internet. Everything from a toaster oven to a nuclear reactor is (or will soon be) connected, and the impact this is having upon security is already overwhelming. These advances are causing a vacuum in IT security management that organizations are rushing to fill.

In this complex threat and vulnerability environment, time is of the essence: The longer it takes to detect events and accurately respond to them, the more damage the organization sustains.

How do we manage risk, given the exploding number of endpoints and the growing threats from those devices and the Internet? How do we get the visibility we need to determine whether a threat appearing to compromise the integrity of an endpoint is real or not? Once we have visibility, how do we minimize our time to respond to such attacks through real-time analysis and remediation? Finally, how do we reduce the time and expense of reporting tasks for compliance or internal audit? This paper explores how to weave together and correlate information from disparately managed systems and bring visibility to their behavior with accurate, actionable reporting in as near-real time as possible—all important components of the Critical Security Controls (CSCs).

1 www.scmagazine.com/suspect-everything-advanced-threats-in-the-network/article/284426

2 www.techrepublic.com/blog/cio-insights/it-security-how-do-you-deal-with-the-unknown-unknowns/39749110

3 www.verizonenterprise.com/DBIR/2013 p. 42. (See “User” category.)

Why Speed and Accuracy Are Crucial

Because the damage can be done in a matter of seconds, we do not protect much, if anything, when we only *discover* an attacker's use of a zero-day exploit through forensic analysis. Consider also that mobile devices can leave the faintest of footprints, often hopping on and off a network in a matter of seconds. Any management system designed to act (or react) to observed events must have as short a reaction time as possible. To avert disaster by blocking a zero-day attack before the attacker takes a strong foothold in the organization, data must be collected, sorted, analyzed and alerted upon within milliseconds.

Putting Intelligence Into Action

For many years now, security practitioners have attempted to address these problems using IDS and correlating data with particular attack scenarios. In this approach, the security practitioner collects as much data as possible from all logging systems, which normalize the data for analysis. Then it is stored in a security information event management (SIEM) system for examination if an attack is suspected. This is certainly beneficial; but in reality, such an approach has us reacting after the attack. A proactive defense requires some sort of "real-time analytics" that can analyze data from the SIEM platform—even better, it includes the evaluation of data as it passes between event sources and the SIEM repository.

Analyzing network activities, system logs, running processes and other real-time data could flag even normal-looking activities. For example:

- A malicious actor uses the Windows Credential Editor⁴ post-exploitation tool to dump password hashes in hope of capturing the hash of a powerful ID, such as an account called *Service Backup*. With real-time analysis, we can be notified when a user or application is logging in with that account outside of the normal backup window.
- Going further, our analytics could trigger alarms and take action to stop the attack, blocking the use of an ID that should not be in use at that time of the day.
- Expanding upon this, we can look for ways to view and analyze our data in real time and react quickly and automatically to protect our systems.

⁴ www.ampliasecurity.com/research/wcefaq.html

Proving It

Consider the legal and regulatory compliance benefits of cross-analyzing data at speeds that can stop attacks. If the system can provide data proving that the attack has been blocked, organizations may avoid more costly remediation of compliance failures. Being able to prove or show examples of protection accomplished before an attack did any reportable damage is much more compelling to compliance officers than are actions taken after the fact.

Earlier this year, Gartner Research Vice President Rita Sallam said in a report quoted by *Computer Weekly*, "Analytics must be made more invisible and transparent, through easy natural language interfaces for exploring data and through embedded analytic applications. Operations staff will increasingly need real time operational BI [business intelligence] systems to be effective."⁵

What can real-time security analytics do to reduce the amount of time corporations must spend to satisfy ever-expanding compliance requirements? Reuters' *Cost of Compliance Survey 2013* notes, "65 percent of firms [surveyed] thought that their liaison with regulators would increase slightly or significantly during 2013. [...] More onerous regulatory and reporting requirements were ... also important factors."⁶ Again, anything real-time analytics can do to improve such reporting will offset its costs, even if only in an intangible sense. Look for reporting that is appropriate to your line of business and its compliance requirements.

⁵ www.computerweekly.com/news/2240177801/Gartner-analytics-will-be-more-transparent-real-time-and-automated

⁶ <http://accelus.thomsonreuters.com/sites/default/files/The-Cost-of-Compliance.pdf>

Applying “Real-Time” to Analytics

Google the phrase “real-time analytics” and the search engine returns close to 9.8 million records.⁷ Most of what that search turned up refers to analysis of business practices to optimize operations and increase profit.

In the security practitioner’s case, we are using the same concept of analytics grinding through unfathomable amounts of data that a human can’t manually interpret. Rather than identify business intelligence, we use security data analytics to identify potential security risk exposure, based on a series of related security events, such as a DNS server making calls to internal systems it shouldn’t be talking to, while said internal system is broadcasting calls to other internal hosts it shouldn’t contact.

Perhaps the most accurate description of real-time security analysis comes from TechTarget:

Real-time analytics ... consists of dynamic analysis and reporting, based on data entered into a system less than one minute before the actual time of use.⁸

In other words, our real-time analytics are not “real time” by the standards of embedded computing, but they are as close to real time as possible. When considering such a system, specific advice for speeding up your analysis includes taking the following actions:

- Centralize device management. Limit the number of consoles and operators who must view those consoles.
- Adapt to multiple devices and data types, normalizing the collected data for further analysis.
- Improve bandwidth and storage to enable speedy acquisition of large volumes of security and vulnerability data.
- Move toward automatic remediation of devices wherever possible.
- Keep “touch points” (points where devices connect) closer together, whether physically, logically or a combination of the two, for faster and more direct communication.
- Recognize wireless speeds are still typically much slower than wired speeds, possibly reducing average collection and reaction speeds.
- Distribute command and communications traffic. Instead of using traditional star or hub-and-spoke topologies, look for products with peer-centric communication, which can provide touch points that are closer to critical devices.
- Tune and configure analysis so that only significant data is processed.
- Only normalize and/or store needed or relevant data. Save only header or metadata information wherever possible when that information is adequate.
- Analyze data for patterns. Compare activity, headers, logs, events and other relevant data.

⁷ As of 10:00am PDT, May 23, 2013.

⁸ <http://searchcrm.techtarget.com/definition/real-time-analytics>

Applying “Real-Time” to Analytics (CONTINUED)

- Ensure that any required agent is minimally intrusive to the host device and has a small memory footprint, low CPU usage, small data packets and no discernable interference with the device or its user.
- Ensure log relevance. Logging should be only as verbose as necessary to minimize impact to productivity.
- Wherever metadata is adequate, analyze that rather than original data to save time and resources.
- Give guidance to users so they know the most efficient way to run queries against the system. If a particular query or syntax is likely to bog down the system, make sure the user understands the implications before he or she runs the query. To preempt such issues, provide a robust, built-in query set that reduces the need for *ad hoc* scripting.

Regardless of the model chosen, the speed at which intelligent parsing of the security- and event-related data occurs is crucial. This may require additional technologies for brokering suspect traffic to inspection devices for further analysis in such a way that good traffic flows without hindrance. This brokering must occur at network speed or better, all while feeding into a larger analysis engine for correlation.

The Importance of Automated Response

Automation is one of five critical tenets of an effective defense based on the Critical Security Controls (CSCs), and one that real-time security analytics directly supports:

Automate defenses so that organizations can achieve reliable, scalable, and continuous measurements of their adherence to the controls and related metrics.⁹

The use of real-time analytics also meets the requirements of multiple individual controls within the CSC guidelines, particularly the need for continuous monitoring.

Automation also adds an incident response angle to this discussion. If continuous real-time analysis ties to appropriate responses, then our analytics system has become an automatic incident response system as well. Shawn Henry, president of CrowdStrike Services and formerly executive assistant director of the FBI's Criminal, Cyber, Response, and Services branch, describes continuous monitoring this way:

... constantly "hunting" for the adversary ... is the most effective way for organizations to reduce the impact from targeted attacks. Assume they're already in. The old performance metric used to be "have you kept the adversary out of the network?"; the new measure today must be "how soon after they make access before we discover them?" That time frame can't be two weeks, two months, or two years. It needs to be two minutes, to minimize the consequences of a breach.¹⁰

Correlation and alerting needs to happen automatically, without the need for humans eyeing every log, packet and bit of security data coming at them—an impossible task given the volume of security data on today's networks. Knowing what to react to and how to react should also be part of an intelligent analytics system. In a recent interview with a SANS editor, a member of the CSO Council of the Bay Area (California) stated that attacks generally take a typical pathway, starting with infection at the endpoint and ending with data exfiltration. Knowledge of this pathway or pattern of events should be part of the SIEM or analytics platform and can help organizations set up "kill zones," to which advanced malware can be diverted—and (hopefully) eliminated.

Many forensics investigators find that the nature of an attack dictates the most effective defense, or the "kill chain," the term used in a Lockheed Martin whitepaper to encompass "the structure of the intrusion, [with] the corresponding model [guiding] analysis [and informing] actionable security intelligence."¹¹ Table 1 shows how threats and countermeasures vary among kill zones.

⁹ www.sans.org/critical-security-controls/guidelines.php

¹⁰ Shawn Henry, *Measuring What Matters In Evaluating Cybersecurity; Social Media Discovers Cyber Security; Bank Shifts Liability for US \$440,000 ACH Fraud to the Victim* www.merit.edu/mail.archives/netsec/msg00502.html

¹¹ *Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains*, 2010; www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf, p. 2

The Importance of Automated Response (CONTINUED)

Kill Zone	Threat	Countermeasures
Endpoints: mobile devices, desktop and notebook PCs, etc.	Users click links or attachments, or are infected by drive-by downloads.	Deploy endpoint monitoring and protections (e.g., antimalware)
Endpoints and servers	Attackers seek out privileged accounts to gain a deeper foothold into the organization, known as <i>privilege escalation</i> .	Access controls and monitoring, including for admin account abuses.
Servers and security	Attackers seek server space to store tools and park data bound for exfiltration. These could stay quiet for a long time.	Monitor communications between servers, particularly DNS servers, security servers and network devices.
Network boundary	Edge servers and routers are often the last point to detect an advanced attack. Criminals know this and tune their outbound traffic to leave the network beneath the radar.	Monitor for short bursts of packets moving out of the network over time. Watch for bursts during hours when users are not usually on the network.

Table 1. Kill Zones, Threats and Countermeasures

One example of an incident that real-time analytics would most effectively counter appears in the final report of the North American Electrical Reliability Cooperative's Cyber Attack Task Force published in May 2012. The hypothetical "Lateral Movements" case study (in the report's Appendix J) describes a malicious actor's careful and patient movements through a network, taking over one system after another.¹²

¹² [www.nerc.com/docs/cip/catf/12-CATF_Final_Report_BOT_clean_Mar_26_2012-Board Accepted 0521.pdf](http://www.nerc.com/docs/cip/catf/12-CATF_Final_Report_BOT_clean_Mar_26_2012-Board_Accepted_0521.pdf), p. 74

What Should Security Analytics Do?

Through analysis of how the network looks when it is healthy, it is possible to program security analytics engines to watch for anomalies and score them based on their propensity for malice. In addition to monitoring through signatures and heuristics, this scoring system can be built based on internal threat intelligence and/or intelligence provided by third parties. Increasingly, forensics firms are marketing “intelligence services,” or “intelligence databases,” while SIEM vendors are either adding additional intelligence analysis into their devices or providing services through partnerships with intelligence vendors, or both.

Key requirements for a real-time security analytics system should include the following:

- Ability to normalize data from an ever-increasing volume of devices, including the latest smart devices, printers and even personal devices (where BYOD is a factor)
- Ability to parse out and identify problems from huge volumes of data quickly and maximize intelligent output from those data
- Automation of analysis and correlations, leading to conclusive probable cause and recommended actions
- Rapid notification of possible incidents sent to appropriate personnel
- Rapid remediation of incidents, preferably before they become attacks
- Ability to easily conduct forensic analysis of incidents
- Continuous tuning of measuring and monitoring thresholds, to keep the system current and trustworthy, thereby minimizing false positives
- Compliance with contractual, legal and regulatory requirements (e.g., HIPAA, PCI DSS, SOX)
- Built-in disaster recovery features that maintain analytics during switchover to/from the recovery environment
- Adherence to relevant standards and best practices—with the capability to accept new data formats, including third-party data

An effective real-time security analytics system should send automated notifications to appropriate personnel as soon as a possible issue is detected. Assuming the system is properly tuned, there will be few false alarms and adequate time to address the alarms that sound. This requires constant tuning, measuring, monitoring and adjustment of thresholds.

Better Analysis Leads to Better Security

By using the Critical Security Controls as a yardstick, real-time analytics should ultimately contribute to measurable reduction of risk. This is accomplished through “rigorous automation and measurement of the Top 20 Controls” as noted by former U.S. State Department CISO, and current director of the National Cybersecurity division of the Department of Homeland Security, John Struefert on the CSC website.¹³ Through awareness of authorized devices and software, as well as regular scanning, patching, configuration and policy updates, such a system will automatically address key CSCs, such as the first four:

- **Critical Control 1:** Inventory of Authorized and Unauthorized Devices
- **Critical Control 2:** Inventory of Authorized and Unauthorized Software
- **Critical Control 3:** Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers
- **Critical Control 4:** Continuous Vulnerability Assessment and Remediation

In addition, real-time analytics can also help with other critical controls, such as defense against malware, or analysis of incidents and appropriate automated responses. Incorporating the critical controls also contributes to compliance with legal, contractual or regulatory requirements that are designed to protect data appropriate to its sensitivity. It is just as important to look for disaster recovery features in a security analytics application. If other operational systems do recover, but security systems do not, then the organization is at increased risk until the security systems designed to protect them do recover.

Combined Intelligence

Zero-day threat intelligence comes from good security analytics. Intelligence is gathered from simultaneously analyzed network traffic between all networked devices, such as servers, routers, mobile and wired computers, and any other devices, whether they are in the official network inventory or not. This traffic is sorted, and key packets and logs are normalized, analyzed, potentially acted upon (based on policy) and stored for further analysis. Additional intelligence into the patterns that advanced threats follow is also integrated into this analysis.

Consider an example involving an unknown mobile device. Note first that it only connects to the corporate network on occasion. Before the device can access the network, a policy enforcement mechanism checks the device for appropriate patch levels, security features, vulnerabilities, forbidden applications and more. This is typically through a device’s mobile device management (MDM) features or with the aid of network access control (NAC) systems. The various systems typically send data from these remediation attempts to a separate management console, and perhaps to the SIEM system as well, where evidence of a zero-day attack may go unnoticed.

Meanwhile MDM and NAC continue to monitor traffic to and from every device, even printers, and treat it appropriately. In this instance, several user applications, access requests and DNS servers are turning up some unfamiliar patterns, including unapproved traffic, failed access attempts and evidence of privilege escalation. A real-time security analytics system should be able to access and consolidate the logs and packets involved in the incident, and take actions accordingly.

¹³ www.sans.org/critical-security-controls

Room for Improvement

Taking action is where most SIEM implementations today fall short. Gartner describes SIEM as supporting detection, incident response and compliance reporting through collection of a variety of event and contextual data sources—the core being collection and correlation of security data from disparate sources.¹⁴ Unfortunately, too many security teams usually treat a SIEM setup as a forensic tool rather than an early warning system. In fact, organizations don't feel that their log and SIEM systems are helping them find advanced threats at all, according to the SANS Eighth Annual 2012 Log and Event Management Survey.¹⁵ This is particularly true in the mobile arena, according to the 2012 SANS Survey on Mobility/BYOD Security Policies and Practices.¹⁶

Today's SIEM platforms need real-time capability or additional integrated intelligence tools to analyze large volumes of data sources and react, repair and remediate events as quickly as possible. The correlation engines should be tunable to cover a variety of potential scenarios, particularly against the attack pathways and footprints of today's attacks.

Ultimately, an organization using real-time analytics should be able to produce measurable remediation results on an automated, 24/7 basis. This, too, is a key tenet of the Critical Security Controls.

Financial Rewards

Cost benefits of real-time security analytics can be achieved on many levels. The first, and of course most obvious, is avoiding costly damage to the organization's reputation and the costs of undertaking cleanup. Another is saving the organization resource time in tracking down and remediating such a disaster. Finally, as less time is needed to find and remediate issues, the costs of impact and remediation go dramatically down.

Say an infected endpoint is attempting to access the network and the attacker tries it three times, succeeding on the third try. Once on the system, the attacker runs a shell to find out what services and applications the mobile endpoint is hosting. Then it waits for the user to request VPN access and hitches a ride into the corporate mail server. From there it's only seconds before the intruder starts searching for administrative accounts, and 10 minutes later, the attacker is running a shell on the mail server. At what point in this process would it be worth money to an organization to shut it down?

14 www.gartner.com/it-glossary/security-information-and-event-management-siem

15 www.sans.org/reading_room/analysts_program/SortingThruNoise.pdf, p. 1

16 www.sans.org/reading_room/analysts_program/SANS-survey-mobility.pdf, p. 1

Better Analysis Leads to Better Security (CONTINUED)

A good explanation of using analytics to reduce attack time by taking the appropriate countermeasures comes from the U. S. Defense Department’s information operations doctrine. In it, an action matrix prepared by Lockheed Martin shows how the phase of an attack determines the countermeasures.¹⁷ The first entry in the matrix related to “Reconnaissance” is *Detect* and the second is *Deny*. Web analytics addresses the need to detect a threat, then firewall and/or block traffic through access controls to deny the attack. If the use of analytics is successful and the attack is defeated at that stage of the action matrix, it removes the need for any other action and illustrates the importance of analytics in addressing persistent threats, as shown in Table 2.

Phase	Detect	Deny	Disrupt	Degrade	Deceive	Destroy
Reconnaissance	Web analytics	Firewall ACL				
Weaponization	NIDS	NIPS				
Delivery	Vigilant user	Proxy filter	In-line AV	Queuing		
Exploitation	HIDS	Patch	DEP			
Installation	HIDS	“chroot” jail	AV			
C2	NIDS	Firewall ACL	NIPS	Tarpit	DNS redirect	
Actions on Objectives	Audit log			Quality of Service	Honeypot	

Table 2. Courses of Action Matrix¹⁸

Although it would be inaccurate to call any benefit from a security system a return on investment, as such, the soft dollar advantage of financial burdens avoided by taking advantage of real-time security analytics has to be factored into any calculations of this sort. We do this by determining the annualized lost expectancy (ALE), as outlined by security guru Bruce Schneier in a 2008 post to his blog, “Schneier on Security”:

Calculate the cost of a security incident in both tangibles like time and money, and intangibles like reputation and competitive advantage. Multiply that by the chance the incident will occur in a year. That tells you how much you should spend to mitigate the risk. So, for example, if your store has a 10 percent chance of getting robbed and the cost of being robbed is \$10,000, then you should spend \$1,000 a year on security. Spend more than that, and you’re wasting money. Spend less than that, and you’re also wasting money.¹⁹

17 www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf

18 Taken from www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf, p. 5

19 www.schneier.com/blog/archives/2008/09/security_roi_1.html

Better Analysis Leads to Better Security (CONTINUED)

The most basic information security courses teach the ALE method of determining how much an organization should be willing to spend to provide adequate security. Although projecting figures that it is impossible to completely validate is a theoretical exercise, in the absence of more concrete numbers, it is the best way to set a security budget. Although simple ALE calculations are closer to guesswork than anything else, they provide a starting point for discussions, one that makes sense for calculating a real financial benefit from any security system.

One alternative to ALE is to use actual case studies as a measure of losses due to IT security breaches. In one high-profile case, executives at Sony indicated that an April 2011 attack exposed data for 100 million customers, shut down the PlayStation network for 23 days and cost the company at least \$171 million.²⁰ If Sony had been able to discover the breach through real-time analytics, before the malicious users had been able to acquire any network data, the damages almost certainly would have been less.

As the proverb says, "An ounce of prevention is worth a pound of cure." What is the value of your organization's data, and how much is avoiding such a breach worth to you and your bosses? No matter how you come up with answers to these questions, you should ask and answer them, not merely as an academic exercise, but as a basic definition of your exposure and risk.

²⁰ www.theregister.co.uk/2011/05/24/sony_playstation_breach_costs

Case Study in Efficiency Improvements

Take the case of hypothetical Large National Bank (LNB), located in Dallas, with more than 100 branch offices located around the state of Texas. The bank has deployed more than 120 servers, some in hardware and others virtually. With 3,300 employees, the bank has more than 3,000 workstations in all locations, with the highest concentration of both at the Dallas headquarters: roughly 1,000 employees and 950 workstations. With a forward-thinking corporate culture, LNB naturally allows its employees to use their own mobile devices, using the now familiar “bring your own device” or BYOD model. In the bank’s estimation, the potential gains in productivity and efficiency offset any additional risk from BYOD.

LNB has acceptable bandwidth among locations, and connectivity is redundant through nearby branches. As a rule, Branch A connects not only with headquarters, but also with Branch B and Branch C. This provides not only a doubly redundant connection to the mothership, but peer-to-peer (P2P) links between locations. Although conventional security wisdom generally frowns on P2P communications, such communications allows LNB’s interbranch messaging to pass quickly through the entire structure without relying exclusively on the central headquarters. This reduces communications latency and allows management tools to quickly traverse the network when pushing patches or querying data.

LNB has set up a single platform to manage multiple security features, applications, and reporting. The platform must integrate with various identity management systems, whether they are based on LDAP, Microsoft Active Directory or something else. The results from improving the system’s efficiency include the following:

- Lower TCO
- Easy administration of devices and applications across the WAN
- Comprehensive visibility of devices on the network
- Faster response to threats or vulnerabilities
- Centralized administration to push agents, updates, patches, and so on
- Automated response and remediation frees up resources for more important objectives
- Reduction of a need for manual intervention
- Inventory of assets across the enterprise
- Compliance with regulations, laws and contractual agreements
- Speedy detection and remediation of policy violations

Harnessing this P2P infrastructure as part of the implementation of real-time analytics will improve the bank’s communications and decrease incident response times throughout the enterprise. It will also make responses easier, particularly to requirements of PCI, SOX and the Gramm-Leach-Bliley Act (GLBA), and to regulatory bodies such as the FDIC and auditing organizations, both internal and external. Faster and easier response translates into a reduced hit on organization resources, and that translates into real dollars on the bottom line.

Special Considerations for the Cloud

Providing this level of analytics and visibility in the cloud is another matter. In 2011, Amazon found widespread misconfigurations and vulnerabilities on the systems of oblivious application owners. In its report, it observed that 98 percent of Windows-based Amazon Machine Images (AMIs) and 58 percent of Linux AMIs contained software with critical vulnerabilities.²¹ On average, 46 services with vulnerabilities were found on individual Windows images, and 11 on Linux images; many of the images were of software that was more than two years old. “Our findings empirically demonstrate that renting and using an AMI without any adequate security assessment poses a real security risk for users,” the authors wrote.

This study shows the need for vigilance through security analytics as applied to cloud-based applications and data. What is the best way to analyze an organization’s infrastructure when it lives in the cloud? How can investigators forensically examine a server that lives in the cloud? These are legal, policy and technical issues that must be considered when developing security analytics in the cloud.

A real-time analytics system must be able to communicate with infrastructure, even if it is virtualized and in the cloud. Monitoring must adapt to this environment, where servers (and some applications) have radically different natures from conventional installations. For example, cloud systems and resources are used on a temporary basis, meaning they cannot be scanned continuously. Worse, from a classic view of best practices, these virtual systems share hardware (and sometimes the same operating systems) with strangers, in the case of public clouds.

Again, the ability to be agile is key, not only in the use of new computing platforms, but also in reacting quickly to issues that arise in those new environments. Tools must collect and analyze data at speeds that might as well be real-time, while reaching into environments that did not exist just a few years ago—without overtly degrading performance.

²¹ www.syssec-project.eu/media/page-media/3/balduzzi-sac12.pdf, pg. 4

Conclusion

Today's faster, newer and greatly improved malicious user requires a faster, newer and greatly improved response. Smarter attackers require more intelligent countermeasures, including faster and more reliable analysis, to react to security issues.

It is no longer sufficient to merely collect security data, chuck it into a SIEM system and analyze it manually. In real life, there is little time to ponder what could be going wrong, what one should do to remediate a potential problem, and perhaps most important of all, how serious the problem may be. An organization without the answers to these questions, lacking visibility into or trust in the accuracy of its security analytics, is more likely to waste resources chasing wills-o'-the-wisp instead of addressing actual dangers.

Speeding up analysis and response through real-time security analytics also provides logical benefits. Again, the primary objective here is to protect the data. If we do that with more efficiency, while requiring fewer resources, that translates to real dollar savings that goes directly to the bottom line. Here, the logical benefits tie directly to the financial benefits and can help offset the mushiness of the reasoning behind the "cost without security" figure an ALE analysis provides.

Using real-time analytics, the security practitioner improves his or awareness of malicious actors and their actions and improves network response and resiliency. Quick analysis of—and faster response to—threats benefit the organization by efficiently protecting valuable data and reputations, as well as providing for follow-up and audit.

About the Author

J. Michael Butler, GCFA CISA GSEC EnCE, is an information security consultant with a leading provider of technical services for the mortgage industry. Butler's responsibilities have included computer forensics, information security policies (aligned to ISO and addressing federal and state disclosure laws), enterprise security incident management planning, internal auditing of information systems and infrastructure, service delivery and distributed systems support. He has also been involved in authoring SANS security training courseware, position papers, articles and blogs.

SANS would like to thank its sponsor:





Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS Riyadh April 2018	Riyadh, SA	Apr 28, 2018 - May 03, 2018	Live Event
SANS SEC460: Enterprise Threat Beta Two	Crystal City, VAUS	Apr 30, 2018 - May 05, 2018	Live Event
Automotive Cybersecurity Summit & Training 2018	Chicago, ILUS	May 01, 2018 - May 08, 2018	Live Event
SANS SEC504 in Thai 2018	Bangkok, TH	May 07, 2018 - May 12, 2018	Live Event
SANS Security West 2018	San Diego, CAUS	May 11, 2018 - May 18, 2018	Live Event
SANS Melbourne 2018	Melbourne, AU	May 14, 2018 - May 26, 2018	Live Event
SANS Northern VA Reston Spring 2018	Reston, VAUS	May 20, 2018 - May 25, 2018	Live Event
SANS Amsterdam May 2018	Amsterdam, NL	May 28, 2018 - Jun 02, 2018	Live Event
SANS Atlanta 2018	Atlanta, GAUS	May 29, 2018 - Jun 03, 2018	Live Event
SANS London June 2018	London, GB	Jun 04, 2018 - Jun 12, 2018	Live Event
SANS Rocky Mountain 2018	Denver, COUS	Jun 04, 2018 - Jun 09, 2018	Live Event
SEC487: Open-Source Intel Beta Two	Denver, COUS	Jun 04, 2018 - Jun 09, 2018	Live Event
DFIR Summit & Training 2018	Austin, TXUS	Jun 07, 2018 - Jun 14, 2018	Live Event
Cloud INsecurity Summit - Washington DC	Crystal City, VAUS	Jun 08, 2018 - Jun 08, 2018	Live Event
SANS Milan June 2018	Milan, IT	Jun 11, 2018 - Jun 16, 2018	Live Event
Cloud INsecurity Summit - Austin	Austin, TXUS	Jun 11, 2018 - Jun 11, 2018	Live Event
SANS ICS Europe Summit and Training 2018	Munich, DE	Jun 18, 2018 - Jun 23, 2018	Live Event
SANS Cyber Defence Japan 2018	Tokyo, JP	Jun 18, 2018 - Jun 30, 2018	Live Event
SANS Philippines 2018	Manila, PH	Jun 18, 2018 - Jun 23, 2018	Live Event
SANS Crystal City 2018	Arlington, VAUS	Jun 18, 2018 - Jun 23, 2018	Live Event
SANS Oslo June 2018	Oslo, NO	Jun 18, 2018 - Jun 23, 2018	Live Event
SANS Paris June 2018	Paris, FR	Jun 25, 2018 - Jun 30, 2018	Live Event
SANS Cyber Defence Canberra 2018	Canberra, AU	Jun 25, 2018 - Jul 07, 2018	Live Event
SANS Minneapolis 2018	Minneapolis, MNUS	Jun 25, 2018 - Jun 30, 2018	Live Event
SANS Vancouver 2018	Vancouver, BCCA	Jun 25, 2018 - Jun 30, 2018	Live Event
SANS London July 2018	London, GB	Jul 02, 2018 - Jul 07, 2018	Live Event
SANS Charlotte 2018	Charlotte, NCUS	Jul 09, 2018 - Jul 14, 2018	Live Event
SANS Cyber Defence Singapore 2018	Singapore, SG	Jul 09, 2018 - Jul 14, 2018	Live Event
SANSFIRE 2018	Washington, DCUS	Jul 14, 2018 - Jul 21, 2018	Live Event
SANS Malaysia 2018	Kuala Lumpur, MY	Jul 16, 2018 - Jul 21, 2018	Live Event
SANS Cyber Defence Bangalore 2018	Bangalore, IN	Jul 16, 2018 - Jul 28, 2018	Live Event
SANS Pen Test Berlin 2018	Berlin, DE	Jul 23, 2018 - Jul 28, 2018	Live Event
SANS Doha 2018	OnlineQA	Apr 28, 2018 - May 03, 2018	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced