

CONTINUOUS MONITORING OF INFORMATION SECURITY

Automated measurement, reporting and alerting can improve the effectiveness of security risk management programs.

Executive Summary

Information security centers around risk management – estimating and measuring risks, defining risk avoidance strategies, controlling and mitigating risks, and reporting on risks. At the end of the risk management cycle is one critical step: monitoring security (hence, monitoring risk). Security monitoring entails examining all of an organization's risk controls, mitigations and policies and answering one key question: Is it collectively effective at managing risk?

Over time, the information security industry has been wrenched back and forth by legislative interest in compliance. Together, the Sarbanes–Oxley Act (SOX) of 2002 and, to a lesser extent, the Health Insurance Portability and Accountability Act (HIPAA) of 1996 have completely reshaped information security monitoring.

Table of Contents

- **2 Continuous Monitoring Components**

- 3 The Value of Continuous Monitoring**

- 4 Implementing Continuous Monitoring**

- 4 What to Monitor**

- 5 Automated Measurement**

- 7 Reporting Tools and Dashboards**

- 8 Alerting and Tracking**

Every security professional's attention has been drawn to something only a few people cared about in the past: compliance. The fallout in the public sector has been no different. The Federal Information Security Management Act of 2002 (FISMA) and the National Institute of Standards (NIST) 800-series security documents essentially define compliance for government.

Compliance has made everyone aware of security monitoring and reporting, but generally from the point of view of an auditor making a point-in-time assessment: Is this organization in compliance today? Have the security controls been effective this month? Did the vulnerability scan find anything this quarter?

Point-in-time security assessments are necessary, but they aren't enough. In addition, savvy information security professionals know they must also have continuous security monitoring. Continuous monitoring is the missing piece to complement point-in-time audits and security assessments.

Continuous Monitoring Components

Continuous monitoring changes the security point of view entirely, yielding a moment-by-moment look into the effectiveness of risk management. It differs from an infinite series of audits performed back to back because it includes three components:

- **AUTOMATED MEASUREMENT** of the effectiveness of security controls and systems on a continuous basis, including as many metrics as possible;
- **REPORTING TOOLS AND DASHBOARDS** that can give both instantaneous and trending information on security status to IT technical staff and management; and
- **ALERTING AND TRACKING TOOLS** that indicate when security controls aren't effective.

The value of continuous monitoring as an integral part of risk management is recognized in those same standards that have outlined compliance strategies.

For example, NIST Special Publication 800-137, *Information Security Continuous Monitoring for Federal Information Systems and Organizations*, explains how continuous monitoring should be implemented as part of the security lifecycle. The Office of Management and Budget (OMB), in Memorandum M-11-33, has made continuous monitoring essentially a requirement for FISMA compliance so that executives can make "credible, risk-based decisions ... on an ongoing basis."

Effective continuous monitoring programs entail more than reading intrusion detection system (IDS), intrusion prevention system (IPS) or data loss prevention (DLP) logs more

frequently. They can fundamentally change the playing field of traditional security processes by shifting security monitoring from a synchronous activity to an asynchronous and reactive activity.

To understand the change in paradigm that continuous monitoring enables, consider this example. A typical enterprise security policy might call for break-in evasion: If someone tries a password three or five or 30 times and gets it wrong, then the account should be locked until someone manually unlocks it.

Mature applications and operating systems support this policy easily. When the monthly or quarterly compliance report is produced, it's also easy to report on who got locked out of which applications and how often.

Continuous monitoring changes the timeframe when a break-in attempt occurs. By watching logs and system status information, a security team knows immediately how many users are being locked out and of which applications.

With intelligent continuous monitoring, an alert can be raised when the rate of break-ins deviates from historical averages. And with reactive continuous monitoring, a user trying to break into an application can be banned not just from that system or that application, but from the entire network.

Active Versus Passive Scanning

Security monitoring entails watching logs, but logs tell only part of the story. Scanners can supplement normal security tools and proactively look for problems, boost visibility and give a head start on dealing with potential issues.

Active vulnerability scanners available from software makers such as eEye Digital Security, McAfee and Tenable Network Security have one piece of the scanning picture. These devices look at systems, probing for services and measuring compliance with patch and version policies. However, active scanners can catch only so much because they have to be told to look for something in particular.

A new technique, passive scanning, closes the gap between active scanning and what is really happening on the network. Passive scanners are similar to intrusion detection system (IDS) units, but they look for something different – unauthorized activity or obfuscated traffic.

Because passive scanners can easily tap traffic anywhere, they can find unauthorized traffic flows that don't pass through normal security devices such as firewalls. Passive scanners also work by using fingerprinting protocols, catching obfuscated traffic that may be operating against policy. As part of continuous monitoring, these devices offer a huge amount of visibility without interfering with traffic.



The Value of Continuous Monitoring

Continuous monitoring moves the information security program and risk management away from a static, compliance-focused view of security to a dynamic view, in which changes in threats or increased risk can receive an immediate response.

Such a program is an obvious requirement in any environment in which threats and risks change rapidly – and it's difficult to imagine any environment that is not seeing rapid changes in its threat landscape.

Continuous monitoring sharpens the focus on what's important. A monthly, weekly or even daily regimen of checking logs and reading reports usually creates data fatigue: The same old information, across all systems and subsystems, soon becomes transparent and unimportant. A continuous monitoring approach brings the important information to the top of the list so that the most significant problems are solved while they are still pressing.

This kind of program brings two main benefits that standard point-in-time security assessments do not: increased visibility and increased control.

INCREASED VISIBILITY enables staff at all levels to see what is happening, as it happens. When information is timely and accurate, everyone in the organization can work together to both understand and mitigate security issues.

Monitoring increases visibility in two ways. First, because security information is being collected all the time, it's easy to present trend information over both short-term and long-term periods. If threat activity is bad today, but more or less the same as it was a year ago, that's one scenario. If the situation is getting worse on a continuous basis, that's a very different scenario, one which requires a significantly different response from both technical and managerial staff.

The second way monitoring increases visibility is by presenting information at the appropriate level of abstraction. At the CIO level, tools such as dashboards help to give up-to-the-minute status information. When management needs to participate, it can immediately see that there is an issue that warrants its attention.

At the technical level, having current and continuously updated information makes staff more effective at debugging issues, understanding potential problems and resolving security breaches as quickly as possible.

INCREASED CONTROL means that information systems and networks can be precisely tuned to the current threat and

risk environment. For many security professionals, system and network configuration and access control are based on a worst-case scenario: If the worst possible thing were to happen, how should security be set up?

Unfortunately, conservative configurations and reactive security approaches stand in the way of getting the job done efficiently and effectively.

When controls are tuned instantaneously to the current environment, the result is greater flexibility at meeting operational requirements. For example, network security tools such as network access control (NAC), when linked with continuous monitoring, can shut down or restrict access to parts of the network by someone detected to be acting "out of profile."

Proactive responses to issues as they occur (applying greater controls to systems and networks), rather than weeks or months later, avoids bigger problems in the future and enables faster operations today.

Security Dashboards

As a part of a continuous monitoring system, security dashboards can provide an at-a-glance view of an organization's security posture. As useful as this may be, presenting an organizationwide view of security in a single screen can be a daunting task.

A good place to start is identifying risk information already present in the enterprise. Key starting points include risk mitigation tools (antimalware, antispam, IPS units), anomaly detection tools (tripwire-style tools, DLP), and network tools (net-flow analyzers and reachability/system status tools).

Analyze each tool's status information to identify measures of security posture and risk. Some tools provide information that's hard to summarize in absolute numbers, which is the most difficult part of building a dashboard.

For each metric, establish a sliding baseline and absolute limits. This makes it easy to determine when any particular metric is out of an acceptable range or norm. Without the context of a baseline, numbers for such things as "viruses blocked per day" are meaningless.

The final step is to create a visual representation that provides a quick snapshot of the organization's security posture. Aim for no more than 12 to 16 panes of data using a color indicator (green/yellow/red is popular) and other easy-to-understand graphic elements, such as dials.

Implementing Continuous Monitoring

The degree of difficulty in introducing continuous monitoring within an organization largely depends on how much monitoring is already happening. For example, if the agency's security team is using an IDS or IPS and has processes in place to review low-priority events and respond to high-priority alerts, continuous monitoring is just a refinement of what is already going on.

However, if the agency is taking a haphazard approach to security information management or depending on a manual review of information, it will find continuous monitoring a bigger challenge. More important, if good security monitoring processes don't fit into an organization for some reason, it probably will not achieve success in implementing continuous monitoring without significant changes in resources, politics or staffing.

If the agency has a FISMA-compliant continuous-monitoring program, it will want to start by consulting NIST Special Report 800-137, *Information Security Continuous Monitoring for Federal Information Systems and Organizations*. While the NIST report is long on organizational advice about how to set up an information security continuous-monitoring program, it stops short of diving into technical details.

For a complete continuous monitoring program, consider how the enterprise will continuously collect and report on nontechnical (management and operational) metrics. For example, many organizations have information security awareness training as part of the general security program. A metric, such as "hours of training per year" or "percentage of staff who have been trained each year," is useful in measuring how effectively this risk management technique is being applied.

When building a continuous monitoring program, the security team must decide whether to integrate continuous monitoring with regular security monitoring or to build a parallel system.

The goal of continuous monitoring is to be aware of whether or not risk management objectives are being met effectively, whereas traditional security monitoring focuses more on reacting to events and resolving issues with users and network-connected devices. The following graph lists the pros and cons to the integration.

Should information security continuous monitoring be integrated with other security monitoring and help desk functions?	
Pros	Cons
Staff are already security focused; the same team and expertise can easily do both.	Traditional security monitoring is generally handled by technical teams, while continuous monitoring must also include management awareness, engagement and decision-making.
Many events will overlap, so handling them with a single team saves duplication of effort.	Troubleshooting and crisis handling tend to consume all available resources, reducing attention to continuous monitoring decision-making.
The technology for continuous monitoring is an extension of what is used for normal security event resolution. Reuse is less expensive and faster than starting from scratch.	Tools for traditional security monitoring may not have the trend and history capabilities needed for continuous monitoring.

In any case, if continuous monitoring is kept separate from existing operations dedicated to security monitoring, the agency's management needs to decide where each type of security event is handled. For example, compliance with patch levels and antimalware updating policies are definitely candidates for inclusion in continuous monitoring.

But if the desktop support team already handles this function, then the continuous monitoring viewpoint on policy compliance is merely one of reporting, rather than alerting or opening trouble tickets. Be prepared for significant amounts of data duplication because both IT operations and continuous security monitoring may generate a lot of information.

What to Monitor

Deciding where to start and what to monitor is the first decision, and it's a difficult one.

Looking at the literature on security monitoring, there's a huge emphasis on the controls: specific requirements for how organizations handle security, such as "limit the number of simultaneous sessions" or "automatically disable account on password failures." The easy view of security monitoring might say, "Let's find a way to measure our compliance with every control, and do so on a continuous basis."

There are several issues to consider when building a continuous monitoring program by focusing on an existing list of controls. This approach essentially duplicates what an auditor is going to look for, which increases the likelihood that the organization pass audits, but it doesn't do anything to further improve the security posture.

Another problem is that this route locks the enterprise into long, static lists of controls (a common federal one, NIST SP 800-53, lists more than 200 controls on 237 pages) that may not have anything to do with how the organization operates or implements security. It also ignores the expertise of security partners by not building on the information available from the many security tools being deployed.

Finally, this approach relies on thinking in terms of the threat environment and security knowledge that existed when the controls were designed, and not what is relevant to today's security posture.

This doesn't mean that control lists are to be ignored completely. In fact, they're invaluable for getting organized about how to monitor security. But in addition to asking, "What should be measured?" the question should be asked, "What can be measured?"

Asking "What can be measured?" means looking at the tools available and leveraging the security expertise of their designers and developers. As the security team explores the capabilities of the products installed in the network and on servers, it will discover many other metrics that are worth measuring – and monitoring – beyond those that appear on published control lists.

There's often a tension between what the agency wants to measure and what can be measured. Dealing with that tension ends up being a valuable contribution to the continuous monitoring program. While there's much to learn from the laundry lists and checklists that others have made, security teams shouldn't be limited by them or hold the organization to metrics that are ill-conceived, repeated by rote, or simply poorly thought out.

At the same time, there's much to learn from security products and processes, and the security team will want to integrate its own knowledge into the monitoring program. Just because a metric from the IDS or DLP platform hasn't made it into an official list of metrics doesn't mean it's not worth adding to the continuous monitoring program.

To move on to the next step of automated measurement, there needs to be a handful of security metrics to measure that will help answer the question "How effectively is risk being managed, right now?" Start with a small set of 10 or 20 metrics, and then revisit this topic once all phases of implementing the initial metrics in the continuous monitoring program have been worked through.

Automated Measurement

If continuous monitoring is to be effective, it must be highly automated. This means that there should be no human action required to generate alerts, review logs or otherwise make a security posture assessment.

Of course, staff must be involved in identifying false positives and in applying proper prioritization to alerts. But staff should be automatically alerted when anything important requires investigation and should be able to verify appropriate security posture every day in less than a minute by consulting appropriate automated indicators.

The goal of all the measurements should be the calculation of metrics – simple numbers that help to measure a particular aspect of security against some allowed limits or a sliding baseline. This means that statistics from much of the data that are available must be generated in order to provide absolute measurements.

A good place to start with automating continuous monitoring is log management. A large part of the information needed in a continuous monitoring program will either come directly from logs, or will be supported by information in the logs. So understanding where logs are sent, and getting good programmatic access to the log system is a logical starting point.

This means that organizations need to create central log repositories for all security-related information. When building continuous monitoring automation, be aware that the security team may have to go to multiple log servers to get the needed information. The following table identifies the logs that are needed (at a minimum) to support a strong continuous monitoring system.

Source of Logging Information	Types of Logging Data and Statistics to be Collected
Firewall traffic logs	Allowed network traffic; blocked attempts
Mail security gateway	Level of mail traffic; counts of malware and spam blocked, as well as unscannable traffic
Intrusion detection/prevention systems	Medium- and high-priority alerts on suspicious and blocked traffic
Network device logs	Switches, load balancers and other devices with SYSLOG capability
Server logs	Windows Event logs from Windows hosts; SYSLOG from Unix hosts
Trouble ticket system	Tickets opened, closed and why

Even with all the logs in one place, the security team may want to take a second look at the deployed security appliances and software to make this task a little easier. For example, consider IDS/IPS deployment.

If a centralized console is being used, access to continuously computed metrics is probably already available, such as top-100 lists (top attackers, attackees, events, classes and categories) that could be difficult to properly regenerate even with a sophisticated security event and information management (SEIM) product with query and alerting capabilities. These top-100 lists make it easy to generate metrics, and they are excellent sources when staff or management want to drill down into the raw data behind a report or alert.

Logs aren't the only must-have piece for continuous monitoring. Vulnerability analyzers and system integrity checkers are also important parts of a security compliance program, and will have useful metrics to include and analyze. While some of what needs to be pulled from logs can be generated by such systems, the level of complexity is high enough that the data may need to be parsed out of reports or special interfaces may be needed to extract important metrics.

The same is true for reachability and capacity measurement systems. As pillars of network and systems reliability, they have a lot of great data hidden in them. With the continuing overlap of network, security and system management, metrics such as "system uptime" and "link utilization" should be part of continuous monitoring systems, even if they don't necessarily map to traditional security controls and metrics.

The following table helps identify the kind of information that should be considered for security metric monitoring.

Source of Data	Types of Security Metric to be Generated
Reachability monitors	System availability data; network latency data
Capacity monitors	Utilization of SANs and disk subsystems; memory of critical systems; CPU levels of critical systems and devices
Bandwidth monitors	Utilization of critical network LAN and WAN links
Vulnerability analyzer	System vulnerability detection; changes in vulnerabilities detected; changes in open ports and systems; time between detection and mitigation
Integrity checkers	Changes in system security settings or registry values; changes to sensitive files or directories

When designing a continuous monitoring program, create at least three tiers to simplify design and maximize flexibility. The top tier should be reporting and alerting engines. These can be small applications that feed data to the agency's existing tools (such as a trouble ticket system or a reporting system), or a larger system with its own alerting and reporting functionality.

The center tier is the monitoring engine that collects statistics, maintains the database for reporting, generates trend information and sends information to the alerting system. A lot of time and money can be saved (and the quality products that already exist in this area can be leveraged) by repurposing commercial or open-source software to be the core of this part of the monitoring system.

At the bottom tier are the actual sources of metrics. Some of these metrics and events will come directly from the end devices in the network (for example, via SNMP polling or traps), but not all. Creating statistics from logs and device web pages or XML interfaces will require some middleware.

Outsourcing Monitoring

Managed-security providers can be a valuable addition to a continuous monitoring project. One of the barriers to successful deployment of security event and information management (SEIM) technology is writing the business rules for the SEIM system.

Getting value out of SEIMs requires business rules that are able to make sense out of log data in the context of the enterprise. No one gets the rules right the first time, which means that experience in particular organizations offers a huge advantage.

When someone knows the organization's operations, it's easier to understand specific concerns from a security standpoint. Managed-security providers can bring that expertise to the table, which pays off in faster deployments and more useful data from the SEIM system.

Another advantage security providers bring to a monitoring project is their multiorganizational view. While there are "big bang" security issues that hit everyone at once, attackers now know that a low-and-slow strategy helps them to get the most value out of their work.

When a service provider can see many organizations at once, it can quickly bring lessons learned from one organization to everyone it serves, offering a better-prioritized response and a higher level of risk mitigation.

When building statistics, avoid the temptation to generate data using scripts running on the center tier. By keeping these functions separate and making a thin middleware layer where needed, greater opportunities arise to use off-the-shelf software for the hard parts of the product, and the agency won't tie itself into a particular product or, worse, a home-grown engine.

The middleware needing to be written will generally consist of a lot of small applications, each written for a specific statistic or metric that's being collected. This might include running queries against the log system, pulling data from web interfaces using XML or screen scraping, or even running shell commands on devices where necessary.

Reporting Tools and Dashboards

Statistics and metrics aren't so useful without some mechanism for reporting them. Therefore, the next tier of a continuous monitoring project should include a variety of reporting systems. In this context, reporting includes both traditional static reports and online ad-hoc or canned reports that are read through a web browser.

It's best to think of reporting in terms of the customers for the reports: organizational managers, technical team leaders covering particular areas and individual technical staff with specific responsibility for a particular security subsystem.

It's not possible to support all these groups with a single set of reports, but constructing multiple report sets with the persona of the report end user in mind will maximize utility and minimize the amount of revision required. Don't be afraid to generate sample reports and ask for feedback early in implementation to be sure the right requirements of each team member are being hit.

As with the monitoring engine tier, it is best to work with off-the-shelf tools, such as reporting systems, rather than invent a whole reporting regimen from scratch. The one exception is a security dashboard (see *Security Dashboards* sidebar), because there is little support in off-the-shelf products for building security dashboards aimed at management users. In that area, the security team is on its own.

For reporting, however, the key strategy is to get data into monitoring engine databases in a format that will make them easy to retrieve with standardized reporting tools. Some of the most useful reporting in this area is time-based reporting, so look for tools that can generate strip charts to show trends over time.

Most people reading reports will not necessarily have the range of acceptable metrics at the top of their head, so appropriate scaling is very important. Be sure to show bands of acceptable, warning and critical values for every graph to make it clear how each metric's performance fits into the big picture of effective risk management.

Reports in printable format (such as Adobe PDF) are the easiest to generate because they don't need to link to supporting data. When making graphs for online review, drill-down capability should be a requirement. At the least, make available different time periods for the same graph (such as daily, weekly, monthly or yearly) and different views of the same graph (such as "only working hours," "24x7," or "only off-hours") to help pinpoint time-based problems.

A stronger approach gives the reader of the report the ability to drill down to specific information supporting the metric. This can include redirects to other management consoles (such as IDS/IPS, vulnerability analyzer, or mail security gateway) or filtered reports from the raw log data that support the score or provide more information.

As the metrics being monitored increase, keep in mind that readers cannot review dozens or hundreds of graphs in each report. Some aggregation of similar values needs to be provided to help reduce the amount of information being consumed.

From there, sub-reports can be generated covering values that are out-of-spec to help draw attention to problem areas needing further review. For example, compliance metrics such as timely application of operating system patches, antimalware updates, personal firewall policy and so on can often be aggregated for a big picture of desktop security.

A strategy combining and separating various metrics will depend on the agency's unique set of priorities and weaknesses. For example, if there is a good track record of keeping servers patched, but a poor record of keeping desktops or notebooks current, then these metrics can be separated to help highlight specific areas for improvement.

On the other hand, if servers and desktops seem to have the same level of compliance, combining them means one less graph that needs to be reviewed. When graphs are being viewed in a web browser, moving from aggregated data to individual graphs should require just a single click.

Alerting and Tracking

Continuous monitoring enables ongoing real-time decision-making when security threats occur. Reporting can be helpful in explaining why there's a problem, but alerting is needed to bring problems to the attention of technical staff and management.

As with all parts of continuous monitoring, the best strategy is to tap existing alerting systems, such as a trouble ticket or help desk system. Recreating such systems from scratch, or installing additional ones, unnecessarily complicates the project.

Continuous monitoring has two main requirements from an alerting system: the capability to open (and close) alerts, and some method of ensuring that alerts are being handled in a timely fashion through escalation and prioritization.

Alerting should be restricted to problems that require action. Any security information that is merely "for information," should appear on daily or weekly reports to keep the stream of alerts as small as possible. While there will always be false positives, especially at first, tuning of the monitoring system and its alerting thresholds should be easy to keep productivity-interrupting false positives (and false negatives) to a minimum, thereby lessening the drain on productivity.

One good strategy is to identify two levels of tolerance for each security metric: a "warning" level that warrants a line in a report or a yellow dot on a dashboard, and a "critical" level that shows up immediately as a red value on the dashboard and generates a trackable alert.

Continuous Monitoring in the Cloud

If good information security and compliance calls for monitoring security controls, then how does this fit in with the move toward public cloud providers?

Cloud service providers are not likely to permit customers to participate in their security monitoring, so any monitoring becomes a contractual issue involving trust, written agreements and the occasional third-party audit.

Cloud service providers, especially upper-layer software as a service (SaaS) providers, come with a greater degree of risk from a security and risk management point of view. The best chance at security success in the cloud is with providers of infrastructure as a service (IaaS). Because services are being used at a low level, the agency can impose compensating controls, such as high levels of encryption of data in transit and at rest.



The information is provided for informational purposes. It is believed to be accurate but could contain errors. CDW does not intend to make any warranties, express or implied, about the products, services, or information that is discussed. CDW®, CDW-G® and The Right Technology. Right Away® are registered trademarks of CDW LLC. PEOPLE WHO GET IT™ is a trademark of CDW LLC.

All other trademarks and registered trademarks are the sole property of their respective owners.

Together we strive for perfection. ISO 9001:2000 certified

108136 – 120302 – ©2012 CDW LLC

