NEXT-GEN UNIFIED SECURITY METRICS

Executive Summary

Cisco has many processes for maintaining security hygiene—patching systems, embedding security, and managing vulnerabilities. Our Unified Security Metrics established actionable feedback loops with IT executives and service owners, leveraging embedded Partner Security Architects to improve closure rates, enable prompt remediation, and improve security performance.

Cisco has evolved those security hygiene processes to integrate measures of operational security risk, and to make those measures visible and actionable for all levels of the enterprise, from the Board down. To provide consistent, repeatable, automated measure of risk we consolidated existing structured data in current systems and leveraged it as proxies for likelihood and impact. The resulting operational security rating (OSR) framework allows clear prioritization, driving visibility and accountability across the organization.

All this has been achievable for little cost beyond that of the integration points for the data, and is providing substantial returns for risk management, automation-driven improvements in timeliness of reporting, and in opportunities to extend the model beyond technical vulnerabilities to other sources of risk such as non- compliance to administrative and operational controls.

This paper discusses the drivers behind the development of the OSR framework, the concepts behind it, the data sources and architecture used, and the outcomes of implementation. It is primarily aimed at information security professionals, IT staff, and business people with an interest in security metrics.

Contents

Abstract

This paper explains how Cisco has evolved its information security hygiene processes to integrate measures of operational security risk, and to make those measures visible and actionable for all levels of the enterprise, from the board down.

Intended Audiences

Security professionals, IT staff, and business people with an interest in security metrics will find this paper of most interest. A basic understanding of security fundamentals is expected.

Introduction

Companies in all industries have been struggling with how to repeatably and consistently track their security posture over time. But the events of recent years have given new urgency to their efforts.

The WannaCry ransomware attack in May 2017 is estimated to have affected more than 200,000 computers across 150 countries, with total damages ranging from hundreds of millions to billions of dollars. No sector of business or society was immune, with essential services unavailable and everyone from companies to individuals finding themselves locked out of their own data.

This was just one of many security incidents that year, some directly impacting the personal data of large numbers of private individuals. What made WannaCry and several other recent breaches so impactful in boardrooms across the land was what they had in common. WannaCry targeted computers through EternalBlue, an exploit in older Windows systems, for which patches had already existed for several months. It was organizations that had not applied these patches that were disproportionately impacted. Similarly, patches for the web-application vulnerabilities exploited in other high-profile incidents had existed, in many cases, months in advance of the actual exploitation of that vulnerability. Nor was this a new trend. To quote Verizon's 2015 Data Breach Investigations Report, "99.9% of the exploited vulnerabilities had been compromised more than a year after the associated CVE was published."

Time and again, organizations that had months in which to protect themselves and their customers failed to do so. The root cause in too many incidents was a failure of basic information security hygiene - not evil geniuses in volcano lairs, but relatively pedestrian attacks with off-the-shelf weapons for which defenses already existed. In such circumstances, the consequences for the victimized organizations can be unusually severe – the failures in due diligence and reasonable care, ordinarily high bars in law, are undeniable in the eyes of the media, the public, and the politicians.

The sight of such major, career-ending, yet avoidable breaches across various industries in 2017 was a sobering one for executives and board members across the country. Many began to ask themselves what issues were lurking in their own information systems that they might not even know about, but which could cripple their companies and ruin their careers overnight.

Cisco: Context, Challenge, and Response

Cisco has of course long had processes for maintaining security hygiene—patching systems, embedding security, and managing vulnerabilities. Through our Unified Security Metrics, we established actionable feedback loops with IT executives and service owners, and worked with embedded Partner Security Architects to improve closure rates, enable prompt remediation, and improve security performance.

In the wake of WannaCry, however, we knew we had to go further. Our new mandate was to enable real-time visibility into Cisco's security challenges, consistently and repeatably identify and surface the most material issues from the Board's perspective, and make those same views available to the executives and service owners across all of our internal clients in order to track and manage those issues to resolution.

To achieve this, we needed to:

- Find a common taxonomy for risk that worked for all internal clients, not just IT,
- Find suitable existing attributes in our data sources that would allow us to assign values to likelihood and impact for the exploitation of a given vulnerability or issue,
- Establish and socialize an Operational Security Rating scale for vulnerabilities, and agree the equivalent mappings for non-technical issues so that they too could be rated and prioritized on the same scale,
- Automate the collection of the data and the application of the Operational Security Ratings per item, and roll them up at application, service, and organization level,
- Last but not least, we had to make the resulting measures visible to the Board, with the same data being made available to the internal clients for incorporation into their own workflows and gates as an integral part of their business operations and decision support.

We were fortunate in that we had begun the journey over a year before WannaCry. We were also fortunate in our partners – InfoSec's Vulnerability Management and Incident Response teams, IT Risk Management and Service Owners, and early adopters on the business side such as Supply Chain and the Customer Experience organizations. All shared a common vision – to understand our risk and manage accordingly – which dovetailed precisely with the Board mandate to make that understanding visible and actionable at every level of the enterprise.

The Cisco Environment

The IT group's core responsibility is to build, deliver, and maintain capabilities to continuously deliver business outcomes with speed, integrity, and simplicity.

Most of Cisco's business processes depend on capabilities delivered by the IT staff, and 99 percent of its data assets are stored, communicated, and processed by the capabilities that the IT group is responsible for.

Cisco IT is focused on services. Everything that the group delivers is a service. Currently there are more than 150 services with more than 2000 applications.

In addition, Cisco uses more than 400 cloud providers.

Every IT service is captured in the service portfolio, every IT application in the application portfolio, with both being hosted in the same database.

Every service has a quarterly risk review with the CIO. Security is discussed there, among other risk measures like compliance with the Sarbanes-Oxley Act (SOX), resiliency, and audits.

IT services have dedicated security primes and partner security architects.

There's also the dedicated InfoSec team, which maintains close alignment with the IT organization but reports up through Cisco's Security & Trust organization to the CSO and COO.

The Security & Trust organization has extended the above model of coverage beyond IT to other internal business clients with their own information systems, such as Supply Chain, Customer Experience, Extranet, and Engineering.

Cisco's mission is to innovate and adopt the most effective information security technologies and policies, share them with our customers, and reflect them in our people, products, and services.

To meet our objectives, InfoSec requires close alignment with other departments. We provide security guidance during the architecture, change management, and operational IT processes.

InfoSec maintains a set of security policies and standards to impose consistency with the IT processes.

Concepts

As already mentioned, through our Unified Security Metrics, we had already established actionable feedback loops with IT executives and service owners to improve closure rates, enable prompt remediation, and improve security performance. The program has had its successes over the years, and established a brand with our IT organization, with USM measures embedded in existing client workflows.

This Unified Security Metric program was therefore our starting point. It was built on the following foundational capabilities:

- Scoring is based on a metric referred to as On Time Closure (OTC). This means that Information system owners are measured against the time to close vulnerabilities in their systems, and the closure of all vulnerabilities are measured against standard SLAs.
- The backend process is semi-automated, which means that still manual tasks and validation steps are needed to ensure the metric data is provided correctly.
- The reporting cadence was Quarterly, and later on became monthly.
- USM measures are vulnerability based, to include Application stack compliance, Design exceptions and Application Vulnerability Assessments (AVA).

In essence what the Unified Security Metric program provided was the ability to measure technical vulnerabilities in information systems, and report on a quarterly basis whether a vulnerability was remediated within the current SLAs.

How did USM measure up against the new mandate? To be successful, we had to:

- Enable real-time visibility into Cisco's security challenges,
- Consistently and repeatably identify and surface the most material issues from the Board's perspective, and
- Make those same views available to the executives and service owners across all of our internal clients in order to track and manage those issues to resolution.

Through analysis, we found three main gaps and recommendations that could lead us to an enhanced unified security metric concept that would meet the new mandate.

- We lacked the ability to prioritize vulnerabilities based on the risk the vulnerability poses to business operations.
- We lacked alignment across existing vulnerability management frameworks within the company, and the ability to 'speak the same language' in a commonly used vulnerability management taxonomy. For example, Cisco as a producer of network and information systems has a PSIRT management process in place, based on CVSS, to rate vulnerabilities found in its own products.
- We lacked real-time vulnerability measures to provide timely visibility to vulnerabilities, and to track their mitigated according to time window SLAs defined by the risk posed.

Therefore our first task was to define our terms around operational security risk in terms of likelihood and impact, both in the abstract and also in terms of existing data attributes already being captured for our information systems, while wherever possible seeking alignment with vulnerability management models on the product side to maximize mutual intelligibility.

Risk and Risk Appetite

Before going into the details of a new risk based concept a basic understanding of 'risk appetite' is needed. For operational security, risk is defined in terms of its potential impact on the ability of the organization to achieve its goals. It follows that the risk appetite of an organization is that degree of risk which its leadership judges acceptable in the pursuit of those goals. That appetite is informed by legal, regulatory, and liability concerns, as well as the balance between the cost of information security systems versus the value protected thereby. Risk can never be managed to zero, and there is no profit without risk. The goal of any Information Security Management System is to keep risk within acceptable boundaries while managing the cost of risk mitigation to justifiable levels.

Key to the success of a vulnerability management and metric solution is to determine the appetite for risk. For a risk based vulnerability management approach, "unacceptable" and "acceptable" risk have to be defined and agreed upon by the organization. Cisco captures its risk appetite in policy, and defines through SLAs the number of days within which any given vulnerability must be remediated for the risk to be considered "acceptable".

The business can also choose to grant an exception for a risk, this means that if a justified business reason weighs out against an unacceptable level of risk, the risk can temporarily be accepted within an agreed upon time frame. Cisco manages exceptions through policy as well, and an exception process is engaged to have an exception granted. All approved exceptions are tracked to closure.

Risk Based Concept: Operational Security Risk (OSR)

To enable vulnerability prioritization, a risk based concept had to be defined, and new measurements introduced.

Because Risk is the key ingredient for the new solution, this concept is referred to as operational security rating (OSR). The OSR is derived from the likelihood of a vulnerability being exploited and the impact an exploit will cause to the organization.

OSR = Likelihood * Impact

In this risk based OSR model each vulnerability is assigned an OSR rating.

R0 – Catastrophic	R1 – Severe	R2 – Major	R3 – Significant	R4 – Moderate	R5 – Insignificant

Each OSR rating denotes a different level of perceived risk, and based on that rating a remediation SLA is defined, on the principle that "Higher risks demand quicker responds and shorter remediation times". There is a difference in enforcement of the SLA between information systems already in production, and information systems in pre-production. Information systems in pre-production are not held against a remediation SLA, but based on the OSR rating a decision is made if the system can go into production or not. The vulnerability management policy enforces the remediation SLAs per OSR-Rating , effectively defining the organization's appetite for operational risk. The table below shows the OSR-Rating definitions:

Operational Security Risk (OSR)	Description
R0 - Critical	ALL HANDS ON DECK - EMERGENCY
R1 – Severe	Danger of material business interruption, damage to our brand and reputation, litigation and/or possible financial loss.
R2 - Major	May result in major damage to our brand and reputation, litigation, and/or possible financial loss.
R3 – Significant	May result in significant damage to our brand and reputation, and/or possible financial loss.
R4 – Moderate	May result in localized harm to our brand and reputation, and/or possible financial loss.
R5 – Insignificant	May result in no harm to brand and reputation, and/or possible financial loss.

With the OSR-Rating defined in the OSR, the next step is to define how Likelihood and Impact are being measured to be able to calculate the OSR-Ratings for each vulnerability.

Likelihood

For operational security, likelihood is commonly defined as the ease of exploit of a vulnerability by a threat. 'Threat' itself has proven difficult to measure, especially for large and complex businesses that operate across the globe. From a feasibility and implementation point of view it is impracticable to automatically and consistently measure internal and external threat factors and weight them in order to derive likelihood. We found the more feasible option was to measure the opportunity to exploit a given vulnerability in terms of the technological sophistication required and the degree to which that vulnerability was exposed to potential adversaries. Therefore we decided to use ease of exploit and degree of exposure in deriving the likelihood measure in the OSR model.

$$\text{Likelihood} = Fn(\text{Ease of exploit, Degree of exposure})$$

For example, an information system hosted in a data center behind several layers of defense and with very limited exposure is far less likely to be compromised than an information system exposed to the Internet, even if they share the same vulnerability.

The ***degree of exposure*** of an information system is defined by its place in the network, and varies from very limited exposure for backend systems behind strictly controlled firewalls to fully exposed systems on a DMZ .

Ease of exploit is defined against the CVSS framework, and also in alignment with other Cisco vulnerability management frameworks, such as Cisco's PSIRT program. CVSS is a published standard, and used by organizations worldwide for capturing characteristics of a vulnerability and producing a numerical scoring reflecting its severity. It also met the key requirement of being immediately available in formats allowing for consistent, repeatable, automated reporting.

Impact

As our proxies for the impact of the exploit of a given vulnerability, we chose two values, the data sensitivity and the information system criticality. In terms of the information security C.I.A triad, the data sensitivity of a system aligns with Confidentiality and Integrity, and system criticality with Availability. As with CVSS, above, these two values possessed the virtue of being immediately available in existing systems of record, allowing for consistent, repeatable, automated reporting.

$$\text{Impact} = Fn(\text{Data sensitivity, system criticality})$$

Data sensitivity is defined by the classification of the data processed or stored by the information system. Cisco uses four data security classes to characterize the sensitivity of its data – Restricted, Highly Confidential, Confidential, and Public.

Application criticality and host priority determine the overall system criticality to Cisco's operation. It is derived from the up time requirements of both the application and the hosts as specified by the business owners. The higher the uptime and availability requirement of the application or host, the more business critical we consider the system to be. At Cisco, application criticalities are classified with C rating, C1 through C5.

Host priority is determined in a similar way: each production host is assigned a support priority level (P1–P6).

Scoring Model

Using data points already captured in existing sources for both impact and likelihood, per system and vulnerability, we are now able to construct an operational security rating scale.

We combined these elements to create the scoring table as shown in the table below. Impact and Likelihood are mapped, and for each vulnerability the relevant OSR Rating is determined.

App Criticality (C) OR App/Device Priority (P) OR Data Classification (R/HC/C/P) **(Impact)**	C1/P1/R	R3	R2	R1	R1	R0
	C2/P2/HC	R3	R2	R2	R1	R0
	C3/P3/C	R4	R3	R2	R1	R1
	C4/P4	R5	R4	R3	R1	R1
	C5/P5–P6/P	R5	R4	R4	R2	R1
		S4	S3	S2	S1	S0
	Severity Rating **(Likelihood)**					

The owner of an application is responsible for remediating vulnerabilities, and once a vulnerability is rated, the owner is notified about the OSR rating of that vulnerability. If a vulnerability is still within its remediation SLA, the risk to Cisco is still considered acceptable. Once it passes that threshold, the vulnerability is classified as OPEN LATE and the risk is considered unacceptable for the organization. The longer that vulnerability remains unaddressed beyond SLA, the greater the risk to the organization, and so that period is tracked as the ***exposure time*** for that vulnerability, and adversely impacts the overall operational security rating for that application and any service relying upon it.

More often than not, any given application has multiple open vulnerabilities. For reporting purposes, when an application has multiple open late vulnerabilities, the OSR rating of the worst open late vulnerability defines the application risk. For each application an 'application score card' is available, and the information on this score card is consumed by those responsible and accountable for remediating these issues on a day to day basis.

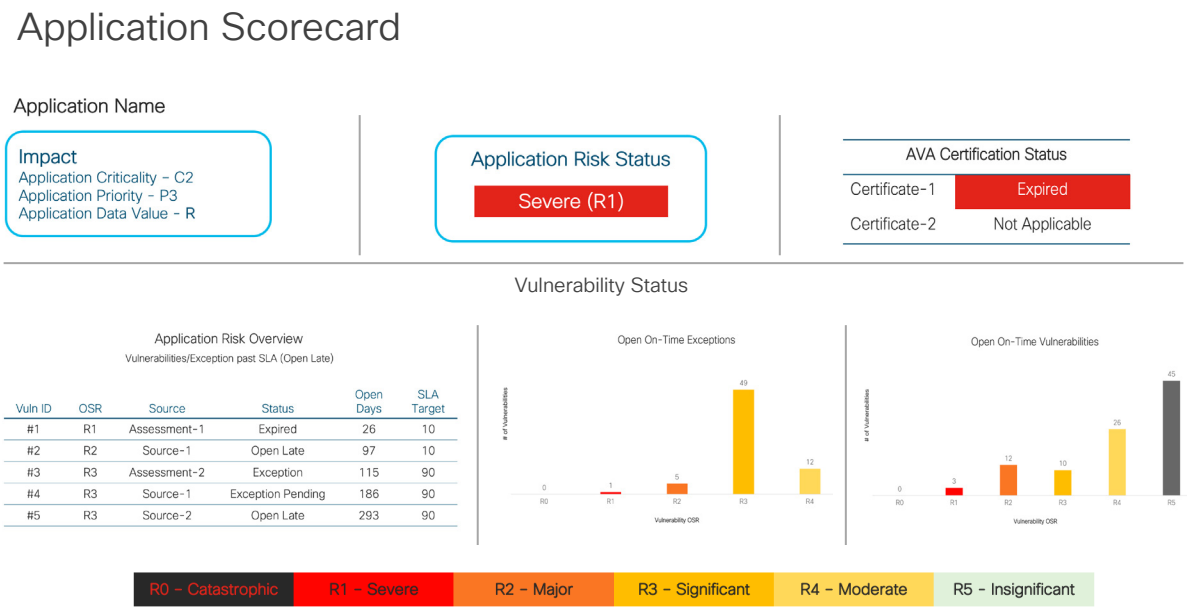

On this example scorecard, the application impact and a top 5 list of open late vulnerabilities are shown. The highest rated vulnerability on the list defines the application risk score. Furthermore an overview is provided for any open on-time pending exceptions and vulnerabilities, providing the system owners with visibility into their OSR 'pipeline'. For each application, the Application Vulnerability Assessment (AVA) status is also listed. Some AVA certifications have to be renewed by application owners on a regular cadence determined by policy. If the application owner fails to comply and the application falls out of certification, this is treated as "out of compliance". In such cases, it is treated as a S2 vulnerability and the OSR Rating is derived on that basis.

The Executive View

Of course, senior executives and board members do not need to be notified about each and every vulnerability. However, executives expect to be informed about the security health of their systems and the vulnerability remediation performance of their teams on a regular basis. Our Unified Security Metrics program already provides a regular cadence and reporting structure for this. As part of an IT Risk Management framework, executives are held accountable for the applications and systems they own on a monthly or quarterly basis. Information security and vulnerability metrics are also part of this framework, and USM leverages this framework to deliver vulnerability

performance and status on a quarterly basis. Alongside the addition of OSR, we have provided executives with the ability to check these metrics on demand, via a dashboard which can provide them with a real time overview of the applications they own.

At Cisco applications are rolled up into services, and executives are responsible for one or more services. Therefore at the executive level all security metrics are rolled up into service-specific overviews.

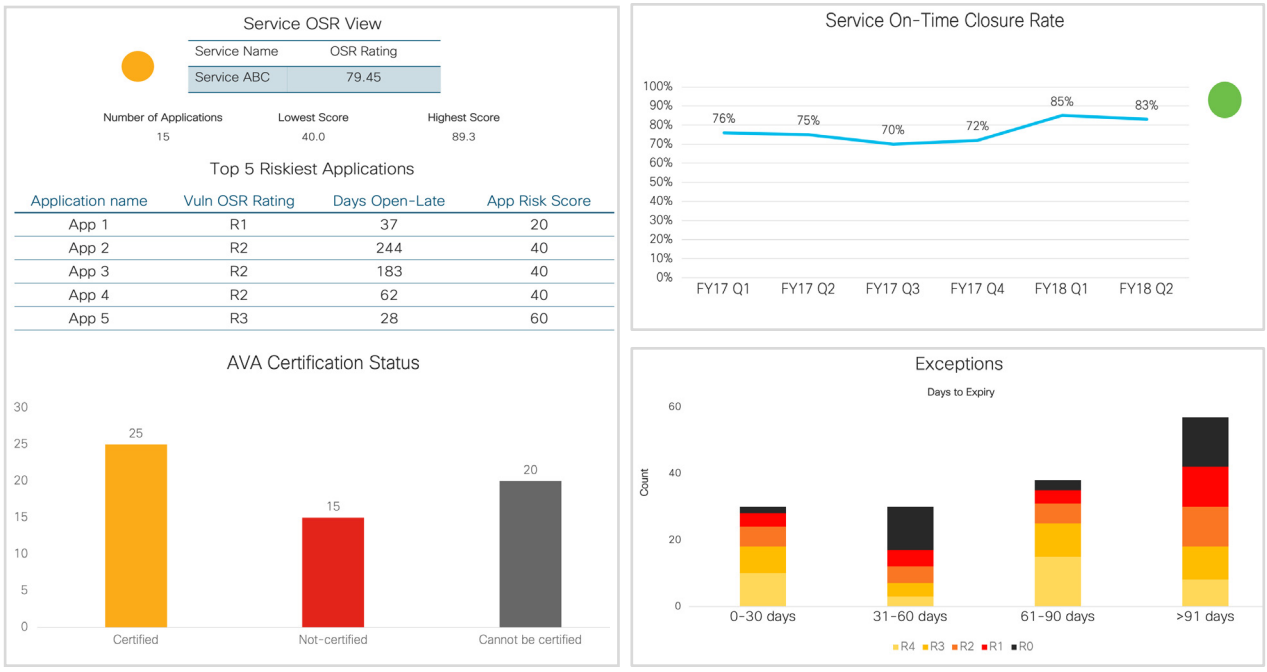

The two main metrics on this scorecard are application risk scores for that service, and the on-time closure rates per risk level. The application risk score informs the executive about the current security posture of the applications delivering the service, whereas the closure rate indicates the degree to which vulnerabilities have been closed on time over the past quarters. For example a service with no currently open R0 and R1 vulnerabilities, but with a low historical closure rate on R0 and R1 vulnerabilities, could indicate that vulnerability remediation have not been properly prioritized in the period prior to the regular date of reporting, resulting in unacceptable risks being taken "while no-one is watching".

The bottom half of the scorecard is used for additional context, AVA certification status of the applications and the number of pending exceptions rolling up under the Services. If this number is increasing it could be an indicator that too many exceptions are being approved to manage vulnerabilities.

Implementation

Sourcing the data

Vulnerability Management at Cisco is a multi-tiered operations. Various teams share the responsibility for identifying the vulnerabilities and remediating them within the SLA. Each of these teams have to adhere to the Information Security policies in this process. The teams maintain their own source of record for the vulnerabilities they identify and/or manage. As a result, the data owners of each vulnerability set maintains control of how frequently the data gets refreshed and the workflow behind each vulnerability that needs remediation.

The tools and assessment processes in place by these teams had previously used different scales for each data set. We have worked with these various teams to adopt a consistent S0-S5 severity scale for these vulnerability measurements. Today **each vulnerability gets assigned a S0-S5 severity level** consistently independent of the source of this vulnerability data.

The **Criticality of the Application, Priority of the Host/Application and Data Classification** records for the applications are maintained in a central data repository. The owners of each application and the teams responsible for managing the hosts are required to register these entities in this central repository. This information can be updated any time as implementation and usage evolves during the life cycle of the application or the hosts. The asset owners are responsible to keep the records updates in the central data repository. This data set too has to meet the guidelines from the Information Security team and the records are governed by internal controls.

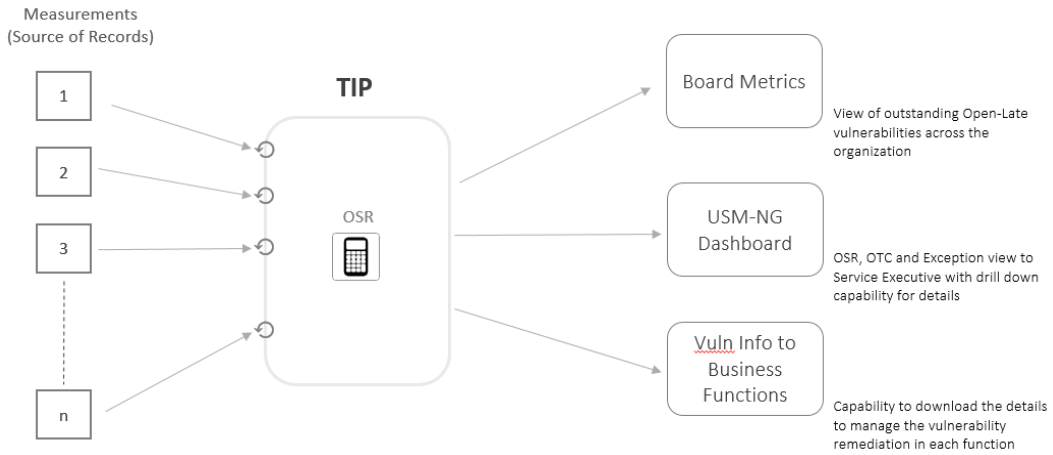

Capability to calculate real-time OSR rating

Combining the data

The measurements we require to calculate OSR score and drive consistency in risk based decisions are in various data sources. Every source of vulnerability information has different needs to maintain these data sets. Based on various factors such as tools used to identify vulnerabilities, assessments that identify vulnerabilities, how frequently the data sets are refreshed and the volume of data generated by the process etc., each team has different tools where these measurements are stored.

To consistently calculate the OSR rating for each vulnerability and maintain a single source of record for OSR rating, we have to ingest these data sets into a central data repository. We have partnered with the big data team, Threat Intelligence Platform (TIP), within Security & Trust Organization (S&TO). The procedure to ingest data from different data sources we use is determined based on the nature of the data set. For example, there are few data sets that are refreshed continuously and we have to ingest these data set every 1-2 hours. On the other end of the spectrum, data sets that are refreshed less frequently are ingested couple of times every day. We have made the determination of how the data needs to be ingested and how frequently we need to ingest the data by working closely with the data owners.

Leveraging the expertise and scale of the big data team has enabled us to keep our metrics always refreshed and reflective of the underlying measurements we use. For example, each vulnerability team reserves the right to change the severity of the vulnerability as they see fit. When the severity level

of a vulnerability changes, we have to accordingly shift the OSR rating of the vulnerability. On the other side of the OSR rating calculation, the type of data processed by an application may get new requirements slapped on it when some data characteristics change. This will again necessitate to change the OSR rating of all the vulnerabilities impacting that application.

It is clear that OSR Rating cannot be a value set in stone for the life cycle of the vulnerability. As the ratings or values of the underlying measurements move, we have to reflect the change by updating he OSR score immediately. This is critical to ensure that all the teams are focusing their effort on fixing the vulnerabilities that pose the highest risk to Cisco.

Moving to real-time

OSR Rating of a vulnerability is calculated in our big data platform every time we refresh one of the underlying measurements. As a result, OSR Rating of a vulnerability will accurately capture the risk to Cisco. This also ensures that we are always reflecting the accurate OSR Rating and remain in alignment to the scoring methodology we have outlined.

Making the data consumable

With real time update of the OSR Rating, we have to ensure that the vulnerability data is made available to the teams in real time to help them manage risk effectively. We have set up a dashboard where each team can track how their service is managing risk. Considering the sensitivity of the dashboard, we ensure that users are authorized to see the services data when they login to the dashboard. A user can only see the vulnerability information for the list of services they are authorized to manage.

Along with the Service views we revived earlier, the dashboard also enables a user to download the underlying vulnerability data for the service. The details will include information about the vulnerability like the source of the vulnerability, the OSR Rating, remediation SLA, status of the vulnerability etc. that help the teams manage the vulnerability remediation process.

To ensure that the user has real time view of the Service risk, the views in the dashboard are generated in real time by connecting to our big data platform. When a user logins in, they will always see the most up to date information about the vulnerability impacting their service along with the capability to download the information.

The risk management programs in different functions at Cisco report various metrics to the Senior Executives in a quarterly basis. To ensure that we make the relevant data available to these risk management programs, we take a quarterly snapshot of the Service views along with the underlying data.

Driving adoption

IT has the most mature risk management program around vulnerability management and application security at Cisco. IT is by far the most robust client of the USM program today and this partnership ensures considerable coverage across Cisco. But there are other business functions that have to maintain their infrastructure and applications outside of central IT environment for various business considerations. We partner with these business functions to drive adoption of USM in their areas.

The OSR score and the associated SLA for the vulnerabilities necessitate a change in how teams manage their vulnerability closure priorities. We have leveraged the prime and partner security architects within IT to drive the change management. The data was made available to this community one quarter ahead of operationalizing this program, enabling them to understand how to access the data, interpret the data and giving them time to ramp up their processes. We had weekly ask anything sessions to give them an overview of the program. This enabled us to clarify the questions the client had along with building a frequently asked questions section that we continuously update.

The journey to adopt OSR and align the remediation SLAs began within S&TO. The first step was to internally build consensus on how we are going to use existing measurements and consistently apply OSR rating to all the vulnerabilities. We worked closely with all corporate functions that manage risk to ensure alignment across all the frameworks already in use. This ensured that USM NG metrics could be reported up the existing reporting mechanisms for visibility and decision support.

As part of the InfoSec outreach efforts, we have security consultants embedded in every business function at Cisco. These consultants wearing a security hat on top of their existing roles ensure that each team complies with the security policies that are relevant to their areas. By partnering with these security consultants, we are embedding OSR- based SLAs in existing workflows. The real time dashboard provides a mechanism for the teams to prioritize the vulnerabilities they should close. Through workshops and S&TO representation at key workflow gates per client we help them understand how they can consume this data for day to day operational decisions.

Lessons learned

The dashboards which leverage OSR enable visibility and drive a virtuous cycle of visibility and action from every part of the organization that is accountable for vulnerability remediation.

Access and common scoring at all levels

Digitized, automated scoring of Operational Security Risk enables us to provide real time status of the riskiest vulnerabilities to the C level executives at Cisco, audit committee and to the Board of Directors. A Board Metrics dashboard showing real time status of R0 and R1 vulnerabilities provides precisely the sort of visibility and accountability that executive suites have been demanding since Wannacry.

At the same time, and using the same data, the USM NG dashboard provides a default view of their vulnerabilities to all the business functions within Cisco. Alongside these default views, we enable each business function to download their vulnerability data with the enriched scoring information, for integration into their own tools and workflows. This empowers each function with the decision support to run their security and vulnerability remediation as they do any other part of their business.

With each function **consuming the data from one central location** we ensure that the metrics we report at **every level** of leadership are **consistent** with each other. Functional leaders know that the data they're seeing is the same that the Board is looking at – a key assurance in driving adoption.

Risk-based prioritization changes behaviors

OSR scoring enables a 360-degree view of the impact from a vulnerability. For example, a host may be supporting multiple applications. When we see a vulnerability on that host, the risk rating for the vulnerability is computed by considering all the impact factors from the applications dependent on the host. This view also enables InfoSec to easily map the risk from this vulnerability if they are having any discussions with the business functions. Enabling this transparent 360-degree view ensures that all the functional teams can work together to remediate the vulnerability within the acceptable SLAs. For example, infrastructure teams who previously found it difficult to arrange remediation downtime with the owners of dependent applications now have OSR data – and the eyes of executive leadership – on their side.

We have an additional benefit of reprioritizing the vulnerability remediation if the likelihood or impact factor in risk score calculation changes. With real time capability, the risk score of the vulnerability can quickly reflect the changes in the underlying factors. This will enable the teams to prioritize or de-prioritize the vulnerability depending on how the underlying factors have shifted.

Authors

Gerwin Tijink, Senior Architect, Information Security

Kaveriappa Muddiyada, Senior Analyst, Security & Trust Operations

Seán Stack, Senior Manager, Security & Trust Operations

Conclusion

The last two years have taught us all some hard lessons about the need for constant vigilance against cybersecurity threats, with some of the worst outcomes occurring through complacency and negligence of basic security hygiene functions.

Maintaining that basic hygiene—patching systems, embedding security, and managing vulnerabilities – is easily overlooked or ignored unless there are actionable feedback loops with IT executives and service owners to highlight issues and risks, drive remediation, and improve security performance.

Cisco's experience with OSR shows that it is possible to provide a consistent, repeatable, and automated measure of risk using existing structured data in current systems. Making that data visible and actionable for all levels of the enterprise, from the Board down now allows clear prioritization, driving visibility and accountability across the whole organization, at little incremental cost beyond that of the integration points for the data.

Working with the grain of the organization, leveraging existing tools and data in new ways, and with patient change management, Cisco has managed to change the way it sees, talks about, and addresses its business risk. Most medium and large companies probably have all the same ingredients at hand – and can use the process outlined in this paper to effectively manage their business risk.