Documenting Cyber Security Incidents

Working paper: Marshall Kuypers (mkuypers@stanford.edu) and Elisabeth Patè-Cornell (mep@stanford.edu)
December 2015

Organizations often record cyber security incidents to track employee workload, satisfy auditors, fulfil reporting requirements, or to analyze cyber risk. While security incident databases are often neglected, they contain invaluable information that can be leveraged to assess the threats, vulnerabilities, and impacts of cyber attacks, providing a detailed view of cyber risk in an organization. This paper emphasizes what data is useful for a risk assessments and how data should be recorded.

When a security incident occurs, it is often unclear what information should be recorded about the incident. Currently, the data that are tracked are largely driven by compliance for reporting requirements and valuable information is not recorded, or information is not recorded in a way that makes analysis easy. For example, incidents are often documented in unstructured reports that require a manual analysis to identify trends. Explicitly recording certain data in structured form makes analysis much more accurate and efficient. This paper is focused on answering the question, 'What should we record?'[1]

Background and other work

Much has been written about data in cyber security. Andrew Jaquith's book, *Security Metrics: Replacing Fear, Uncertainty, and Doubt*, is a good primer on data collection and includes a list of data sources, computational techniques, data visualization methods, and many discussions of the usefulness and limitations of metrics [1]. In 2010, the Center for Internet Security organized 150 experts to create a list of data that should be recorded, and what metrics should be reported [2]. NIST special publication 55, revision 1 is a guide to measurements in information security [3]. Other research focuses on the processing and aggregation of data from sources like intrusion detection systems and other log files [4]. Recently, interest in dashboarding has increased, and significant literature exists on data visualization[2].

Several frameworks exist for recording information about cyber security incidents. For many years, the US CERT Federal Incident Reporting Guidelines mandated what information was collected and how it was reported for certain organizations. The general information reported included the incident date and time, attacker/ victim IP, port/ protocol information, operating system information, detection information, impact, resolution, and incident category. The incident categories are listed in the table below [5].

[1] Other work emphasizes what metrics should be reported in a weekly summary (e.g. number of unpatched systems, mean time between incidents) to give the best situational awareness. Many of these metrics are derived from the information that is recorded, while others are measurements that describe an organization's current state (for example, the number of unpatched systems on the network). This paper does not emphasize these questions, but instead focuses on what data about a cyber security incident should be recorded. The implementation of an effective process that identifies, categorizes, and documents incidents is also nontrivial, but not the focus here.

[2] See Edward Tufte.

Federal Agency Incident Categories

Category	Name	Description
CAT 0	Exercise /Network Defense Testing	This category is used during state, federal, national, international exercises and approved activity testing of internal/external network defenses or responses.
CAT 1	Unauthorized Access	In this category an individual gains logical or physical access without permission to a federal agency network, system, application, data, or other resource
CAT 2	Denial of Service (DoS)	An attack that *successfully* prevents or impairs the normal authorized functionality of networks, systems or applications by exhausting resources. This activity includes being the victim or participating in the DoS.
CAT 3	Malicious Code	*Successful* installation of malicious software (e.g., virus, worm, Trojan horse, or other code-based malicious entity) that infects an operating system or application. Agencies are NOT required to report malicious logic that has been *successfully quarantined* by antivirus (AV) software.
CAT 4	Improper Usage	A person violates acceptable computing use policies.
CAT 5	Scans/Probes/ Attempted Access	This category includes any activity that seeks to access or identify a federal agency computer, open ports, protocols, service, or any combination for later exploit. This activity does not directly result in a compromise or denial of service.
CAT 6	Investigation	*Unconfirmed* incidents that are potentially malicious or anomalous activity deemed by the reporting entity to warrant further review.

In 2014, in an effort improve data quality, improve compatibility with other guidelines, and to increase the speed of information sharing, US-CERT revised the reporting guidelines with an implementation deadline of September 2015 [6]. The biggest difference in the new reporting guidelines includes a qualitative assessment of the impact and the identification of the threat vector. Unfortunately, the inclusion of qualitative (low, medium, high) assessments of impact have minimal value compared to quantitative estimates (for example, investigation time). Further, it is uncertain if the threat vectors are of sufficient clarity[3].

[3] For example, vectors and exploits are combined. It is unclear if a brute force attack against an exposed administrator login on a website would be classified as ‘Attrition’ or ‘Web’. Attrition is a technique, not necessarily a vector. Similarly, impersonation is a technique, but could be used through an email vector (phishing), a website vector (SQL injection), or a physical access vector (delivery man costume).

Other frameworks for recording and reporting incidents exist as well. VERIS (Vocabulary for Event Recording and Incident Sharing) is a well-documented framework for reporting incidents[4]. VERIS structures information for incident tracking, victim demographics, an incident description, the discovery and response, and an impact assessment. One of the drawbacks to VERIS is the large number of fields, which can create data fatigue for investigators repeatedly entering the same information for low-impact incidents. Other taxonomies can be useful to compare incident records against. MITRE has a characterization for specific attacks called CAPEC (Common Attack Pattern Enumeration and Classification)[5].

Recently, a working group at the Department of Homeland Security has published guidelines for incident recording via an excellent whitepaper [7]. Additionally, several companies have recently entered the incident management software market, suggesting that incident data will become more prevalent in the future.[6]

What data is Useful?

Many organizations record cyber security incidents in some form of database, ranging from ticketing systems, to excel spreadsheets, to home-built software solutions. The cyber security incident databases will contain incidents that are investigated by a security operations center investigator[7]. Each incident will have information about when it was created and closed, who the investigator was, incident information (vector, machine type, etc.), and impact information (hours of investigation, monetary costs, etc.). The incidents will often contain both structured entry fields (number of hours, data encrypted? yes/no) and unstructured text descriptions[8] of the incident.

Organizations accumulate data at different speeds, depending on their size, attack surface, and defenses. An organization of 5,000 employees may only record hundreds of incidents per year. Still, after a short period of time, enough data can be gathered to get a first-order approximation of cyber risk.

Emphasis should be placed on recording data that is most relevant to the cyber security questions at an organization, and in a way that makes analysis simple. Careful thought should be put into the following:

What does the recorded data show?
How does the data lead to actionable decisions?

For example, recording how many websites are compromised via SQL injections versus cross-site-scripting attacks is not important unless a safeguard has a different effectiveness for one vector over the other. Maps of attacker IP addresses are interesting to look at, but often are indistinguishable from population

[4] http://veriscommunity.net/index.html

[5] http://capec.mitre.org/

[6] RSA Archer offers incident recording software, along with Resilient Systems, CyberSponse, and Demisto.

[7] Note the distinction between log data and security incident data. Log data is machine generated and security incident data is generated by a security investigator and summarizes a security incident.

[8] A common challenge is that the structured fields do not have the flexibility to capture sufficient incident detail, so the actual incident description is largely captured in the unstructured text section. This makes analysis particularly difficult, because analysts need to use natural language processing or a manual analysis to extract useful information from the database.

maps and offer very little actionable intelligence. Still, recording more information that may become useful at a future date is preferable to ending up with critical information missing.

Subjective and uncertain assessments can also be valuable. Recording the suspected adversary (nation state, activist, or spammer) is difficult and uncertain, but can lead to richer conclusions than simply collecting attacker IPs, which can easily be masked. Security operation centers usually have decades of combined experience and are skilled as using intuition for characterizing types of attacks. Certain malware variants will be routine, while others attacks can be identified as more serious based on the overall incident. The value of these subjective assessments should not be ignored.

Data Collection Examples

Field	Variables	Comments
Device Type	Mobile device, laptop, desktop, server, token, other	At one organization, stolen and lost devices were rising over time. After looking into the data, it was determined that the rate of lost laptops was constant over time, but more security tokens were going missing. The increase in lost tokens was well explained by the number of employees who had security tokens. As the organization had rolled out the new technology, they would expect that the number of lost devices would rise as well.
Sensitive Information Compromised	SSN, email addresses, salary information, classified information, trade secrets, etc.	Identifying the type of information compromised helps organizations identify attacks with the greatest risk.
Level of Access	None, application, local, root	Acts as a proxy for incident severity.
Alert Source	3rd party, system admin, IDS, other security application, user	An organization analyzed how their security operations group detected malicious emails that were sent to a user. They found less than 15% of users reported malicious emails that they received. This information should motivate the CISO to reach out to the workforce to inform them that sending malicious email to abuse@company.com is a valued use of their time.
Suspected Adversary	Nation state, Lone Hacker, Criminal Org, Spammer, Ideological hacker, Street criminal, Employee (misuse)	Attribution of cyber incidents is difficult and often, too little information exists to make a good determination (for example if a system is scanned). However, the security operations investigators have significant expertise and can often assess a suspected adversary based on the type of attack. For example, common malware exploits can often be attributed to lone hackers or criminal organizations, while specialized malware is clearly the work of more advanced attackers.
Hours of investigation	Numerical	Hours of investigation is a significant cost for organizations. Explicitly recording the hours gives great information to analysts about the resources needed in different areas.
Other costs	Hardware replacement, credit monitoring, loss of revenue from downtime, 3rd party investigation costs	Recording other financial costs allows organizations to connect incidents to monetary losses more effectively. The majority of incidents will not have other costs associated with them, but identifying the few that do is very valuable.

The value of Tags

A major challenge in incident data recording is collecting the information in a way that makes analysis easy. Too often, the data is in a format that either offers too little detail to be interesting, or too much detail (in the form of unstructured text) so that a manual analysis is required to derive useful insights. Introducing incident tags on the fly can alleviate data analysis challenges.

Tagging incidents involves entering keywords to associate incidents with different categories. For example, #lost would denote devices that are lost or stolen. Without tags, identifying a basic incident like lost or stolen laptops can be difficult. Lost laptops can have several words associated with them, including lost, loss, stolen, missing, and found. A search for these words in unstructured text will return many false positives, making the problem labor intensive. Creating an unambiguous tag that is written whenever a device is lost makes the analysis much easier.

Multiple tags can be used for an individual incident and tags can be created whenever new incident types occur. For example, #heartbleed and #shellshock can quickly be associated with new attacks. Other subtle attack patterns can be identified as well. An investigator might begin to notice that many of the website defacements that they are dealing with were created by interns, but never transferred to a full time employee, #abandonedWebsite[9]. Creating a tag for this scenario begins to track how big the problem is, and can track the effectiveness of new policies as they are rolled out.

Tags enable trends to be visible. Data spillage is an issue at many organizations[10]. Without appropriate tags, data spillage incidents can be scattered across many categories (websites, email, physical access, etc.). Creating a flag makes the issue traceable. Similarly, false alerts, near-misses, and non-incidents are often lumped together. False alerts are incidents that do not correspond to a potential threat. For example, a website might be flagged as malicious because it was recently created for an academic conference, and had not been scanned for malicious content[11]. A near-miss is an incident that could potentially cause harm, but did not have a negative outcome due to chance. For example, a Windows virus downloaded on an Apple product is a near-miss. Non-incidents are not security incidents and have been created in error. For example, some security incident database management systems do not have the ability to delete incidents, so if two investigators open a ticket without communicating, one of them will be closed as a non-incident. There is significant value in distinguishing between these incident types. Near-misses are valuable learning experiences, while non-incidents should immediately be eliminated from the analysis.

Where to find Data

Many organizations that do not have security incident databases still have useful data in other locations. IT procurement likely records laptop thefts due to the need to order new devices, system administrators have log data, HR should have good visibility into personally identifiable information, and the legal

[9] Note that the tag label is irrelevant. Emphasis should be placed on ensuring analysts and investigators understand its meaning, which may be refined over time.

[10] Data spillage is the unauthorized and accidental release of information. For example, an employee forgets to encrypt a sensitive email, a tax form is sent to the wrong recipient, sensitive documents are left in the printer room, or a classified document is accidentally posted to a public website.

[11] Another false alert is an email flagged as spam due to poor grammar.

department should records on data breach disclosures and intellectual property loss. Jaquith's book gives an excellent overview of other sources of data including asset management systems, patch management systems, and audits. Even organizations without mature cyber security incident tracking have significant data sources; they just need to be identified. Other sources of data are also available for assessing cyber risk. Industry reports, academic papers, and subject matter experts can all be useful data points.

Conclusion

Incident recording will never be perfect. Cyber security is a rapidly evolving field and new attacks will emerge. Fielding a flexible incident recording framework allows organizations to quickly identify new trends, improve situational awareness, and assess cyber risk effectively.

List of data to record (via structured fields or tags)

General information
Date/Time incident opened
Date/Time incident closed

Impacts
Hours of investigation
Downtime to employees
Credit monitoring costs
Hardware replacement costs
Other costs

Lost/ Stolen devices
Stolen device
Lost device
Device type
Device encrypted
Device shut down or in sleep mode
Sensitive information on device

Websites
Website defacement
Website database dump
Website compromise
Lateral movement
Website on DMZ
Website patched
Website actively managed

Email
Number of targets in organization
Targeted to individual
Spammer
Malicious link
Malicious attachment

Malware
Malware variant
Web browsing
Appropriate website
Outside institutional firewall
Potentially unwanted program
Infected USB

Other
Detection source
DoS
Data spillage
Suspected adversary
Targeted to organization
Brute force

References

[1] Jaquith, Andrew. Security metrics. Pearson Education, 2007.

[2] The Center for Internet Security. "The CIS Security Metrics" November 1, 2010. https://benchmarks.cisecurity.org/tools2/metrics/CIS_Security_Metrics_v1.1.0.pdf

[3] Chew, E., Swanson, M., Stine, K., Bartol, N., Brown, A., Robinson, W. Performance Measurement Guide for Information Security. National Institute of Standards and Technology Special Publication 800-55 rev. 1. July 2008. http://csrc.nist.gov/publications/nistpubs/800-55-Rev1/SP800-55-rev1.pdf

[4] Vaarandi, R., & Pihelgas, M. (2014, October). Using Security Logs for Collecting and Reporting Technical Security Metrics. In Military Communications Conference (MILCOM), 2014 IEEE (pp. 294-299). IEEE. https://ccdcoe.org/sites/default/files/multimedia/pdf/milcom14-metrics-web.pdf

[5] Federal Incident Reporting Guidelines. US-CERT. https://www.us-cert.gov/government-users/reporting-requirements

[6] US-CERT Federal Incident Notification Guidelines. US-CERT. https://www.us-cert.gov/incident-notification-guidelines#Standard_Data_Elements

[7] “Enhancing Resilience through cyber incident data sharing and analysis,” DHS whitepaper. September 2015.