Global Information Assurance Certification Paper

Incident Tracking In The Enterprise

GIAC (GCIH) Gold Certification

Author: Justin Hall, justin.hall@cbts.net
Advisor: Dr. Kees Leune
Accepted: June 28, 2015

Template Version September 2014

Abstract

Some organizations employ Computer Security Incident Response Teams (CSIRTs) to investigate and respond to security incidents. They often find these investigations to be poorly executed, time consuming, and ultimately ineffective at discovering the root cause of a breach. Unfortunately, this is not usually due to the skill of the investigators, but rather due to the tools and processes they use to manage the investigations. This paper describes the use of purpose built case management software, integrated into the incident response process, to track these investigations. CSIRTs that take an organized, formal tracking approach will collaborate better and find their investigations to be more complete and useful to risk managers.

1. Introduction

Computing environments are frequently subject to adverse events, which are 'events with a negative consequence such as system crashes, packet floods, unauthorized use of system privileges, unauthorized access to sensitive data, and execution of malware that destroys data' (Cichonski, Millar, Grace & Scarfone, 2012, p. 6). In recent years, these adverse events have been less directly destructive and included theft of sensitive data or embarrassment of an individual or company. Adverse events sometimes violate 'computer security policies, acceptable use policies, or standard security practices'; and in these cases, they can be referred to as *computer security incidents* (Cichonski et al, 2012, p. 6).

To help protect their assets, organizations often employ a Computer Security Incident Response Team (CSIRT), a group whose responsibility is to detect and respond to computer security incidents. Responding to these incidents often requires a formal investigation, where the CSIRT will gather and analyze evidence in an attempt to fully understand the details of the incident.

The process of gathering and analyzing evidence from workstations, servers, mobile devices, network infrastructure, and users is often time consuming and resource-intensive. It requires management of many complex pieces of information. Examination of a single malicious program can often generate dozens or even hundreds of pages of data, which can be used to identify other compromised computers, or detect additional attacks by the intruder. This data can include Internet protocol (IP) addresses, other

Justin Hall, justin.hall@cbts.net

computer names, forensic artifacts left on the hard drive of the compromised computer, and pieces of computer code.

Investigators must possess sound organizational and analytical skills to properly catalogue and apply this information. The nature of the information - data about the operation of computer systems and networks - suggests that computer systems should be used to store, analyze and share it. In addition, the process of investigation itself is best automated using computer software, to ensure completeness of the investigation and consistency between incidents and CSIRT members. Collectively, managing a computer security incident investigation and the volume of evidence related to it can be referred to as *incident tracking*.

2. Effective incident tracking

This paper recognizes that many factors come into play when assessing the effectiveness of a CSIRT or its investigators. The goal of the paper is to focus specifically on the usage of computer software and systems to store, analyze and share incident tracking data. This includes both the data gathered and created during the investigation, as well as the workflow of the investigation itself.

2.1. Benefits of effective incident tracking

Effective incident tracking, then, involves competent application of appropriate computer software and systems to manage investigations. There are many benefits to this practice.

2.1.1. A More Complete Investigation

Investigators using incident tracking tools to manage the incident response process should be more thorough in their investigation practices. These tools can be customized to employ workflows, ensuring that the correct team members in each

investigation perform consistent steps. The tools can also require specific pieces of information to be gathered before an incident is closed or moved to the next workflow step. Without these requirements, investigators may forget to or choose not to perform certain investigatory actions or collect certain data and miss key evidence or make poor analytical decisions. With them, the CSIRT can be assured that all incidents are investigated using the same, repeatable processes that have been decided upon by the team.

The potential exists for evidence or notes to be lost during an investigation if an investigator stores their data locally on a forensic workstation or in personal records and does not handle that data responsibly, or inadvertently loses them due to hardware failure. Incident tracking tools typically run from a centralized location and store the entire team's data in a database, which can be backed up according to the organization's policies and in a more effective and consistent manner than an individual might use.

2.1.2. Better Collaboration

Most CSIRTs are made up of several team members, often focused on different areas of the computing environment or with different expertise. Sharing data between multiple team members, using document files, such as spreadsheets or word processor files, and utilizing messaging tools like email presents challenges. Rarely do all team members possess and work on the latest data, and it is difficult to effectively aggregate all data into one individual file that can be accessed and updated by all team members at once.

Incident tracking systems, which are typically ticket- or case-based, can handle the management of multiple users updating a master incident ticket. Team members can access the master ticket without worrying about whether the data in it is fresh or stale. Data collected by each team member can be submitted, accessed and searched by all from

a single interface with no need to manage network file shares or dig through cumbersome email archives.

Using workflows in the incident tracking tool can improve parallelization of tasks. From one location, data collection, analysis, development of new intelligence, and remediation tasks can be assigned to multiple users at once by the incident response lead, and all can carry out their tasks and report back to the master incident ticket.

This is especially useful when CSIRTs need to leverage IT operations teams, such as network administrators or help desk staff, to carry out tasks such as implementing firewall rules or reimaging a workstation. Rather than having to transfer data manually from the CSIRT to the tools used by an IT operations team, a task can be assigned from within the incident tracking tool, which will typically have a way to set up one- or two-way interaction with IT operations' ticketing systems. Some CSIRTs provide access to the incident tracking system directly for relevant IT operations staff. This way, the IT operations staff can always see up-to-date information and status from the source.

2.1.3. Historical Tracking & Threat Intelligence

A good CSIRT will leverage data from past investigations when beginning a new investigation - determining if a tool, credential or host has been involved in a previous incident. This can save valuable analysis time and help establish a threat's *modus operandi* - its common habits, strategies or patterns. This kind of searching can be difficult when using individual files stored in many locations and in disparate formats.

Using historical data is typically trivial when using an incident tracking tool. An investigator can quickly search all data submitted by various team members and benefit from that analysis. This may be as simple as searching for a specific domain name or IP address to find other threats or campaigns that have used them; this can save the investigator time in determining which tools might be used, or what the attacker's goal

might be during the most recent incident. At times, new data discovered during an investigation might even cause a CSIRT to examine previous related incidents and reopen cases using the new information.

This data is often valuable to a threat intelligence team - one devoted to developing and implementing indicators that a threat is active in the environment, which can be used by security monitoring tools. These tools, including intrusion detections systems (IDS) and security information and event management (SIEM) tools, need continuous updates to signatures or behavior-based rules to effectively detect suspicious activity.

Typically, these indicators come from the vendor of the tool - the developers of an IDS will release new sets of signatures that are useful to their entire customer base, which detect a wide range of exploits. But, these signatures may or may not match the activity of an intruder that has actually targeted the organization. More valuable to the organization are indicators that reflect the attacker's actual tactics, techniques and procedures (TTPs).

A threat intelligence team is tasked with analyzing the evidence left behind from successful and failed attacks and developing countermeasures to detect and, often, deter future attacks. To this team, nothing is more useful than real data about actual attacks, especially those that have been thoroughly investigated, and whose key indicators have been documented in a continuously updated database. The incident tracking system can serve this purpose.

Threat intelligence analysts can regularly review new incident tickets - both for successful attacks as well as false positives - and use the discovered information about threat TTPs to improve their signatures and rules. For example, it would be evident from incident tickets describing the activity of the threat actor if a command and control host

had not been used by a specific threat in some time, or a threat actor has clearly migrated to a new one. Or, if an updated version of a backdoor has been used in a new attack, and an investigator discovers it, the sample, or analysis report for the sample, could be attached to an incident ticket and reviewed by the threat intelligence team.

Some threat intelligence teams may even configure their intelligence database application, such as MITRE's *CRITs* tool, to automatically pull certain data feeds from an incident tracking system, so that they can rapidly deploy new indicators that have been uncovered during incident investigations.

2.1.4. Improved Reporting

Risk managers often lament the lack of real data available from which to make risk decisions. Darril Gibson (2011, p. 6), in his book on information systems risk management, says that '...previous security incidents are excellent sources of data. As evidence of risks, which already occurred, they help justify controls. They show the problems that have occurred and can show trends'. However, Gibson (2011, p. 20) also acknowledges that 'ideally, weaknesses from a security incident will be resolved right after the incident. In practice, employees are sometimes eager to put the incident behind them and forget it as soon as possible'.

Maintaining a persistent database of detailed incident information would help to solve this problem for managers. Conducting searches or running automated reports across all incident tickets would allow views into targeted data sources and records, applications and platforms, and geographic locations and users. It would provide insight into which attack vectors and tools are of highest concern, as well as the controls and defenses that are most effective at stopping or slowing these attacks. Finally, the technologies that can provide CSIRTs and security groups with the most visibility, through logs and application programming interfaces (APIs), can be identified.

Justin Hall, justin.hall@cbts.net

With this real data, risk managers can make more informed decisions. This situation may also end up benefiting CSIRTs and security teams as well - managers may decide that these teams have inadequate resources and choose to bolster them with additional staff, tools, or training.

2.2. Best practices for effective incident tracking

Effective incident tracking should follow a formal process, involving the appropriate staff and capturing certain pieces of data during each case. CSIRTs should consider how data from individual incidents could be linked together so that patterns and trends in the data can contribute to both the incident response mission as well as the overall security posture of the organization.

2.2.1. Build IR process into tracking tools

Mature CSIRTs typically separate the incident response process into phases and use these phases to guide the work of investigators, analysts and IT operations staff. SANS recommends a six-phase process as a part of their SEC504 course - preparation, identification, containment, eradication, recovery, and lessons learned (Skoudis, 2010, p. 14). Similarly, NIST recommends preparation, detection & analysis, containment/eradication/recovery and post-incident activity (Cichonski et al, p. 21). Each of these phases typically involves several tasks, many of which are dependent on the outcome of a previous task.

Each task may involve the work of a specific team member or one of many that fills a particular role. During the identification phase, for example, an alert may need to be triaged, which means its validity and severity must be determined and a course of action decided upon. This may be the initial step in the incident response process. Following triage, a confirmed compromise may be sent to an analysis team to determine the extent of the attacker's activity and presence on the network; in the case of a false

positive, the initial alert may be sent to a security operations or threat intelligence team to find the root cause and update signatures in use on detection systems.

It is usually possible to construct these kinds of workflows within the incident tracking tools. The incident and its relevant data are stored in a single unit, usually a ticket. Investigators become task owners, and can be created as individual users or groups. Teams can be created that handle specific phases, and queues - lists of pending or actionable tasks or tickets - can be set up so that these teams need only to review them to know which items need their attention.

Finally, the progression of the investigation through the IR process can be designed as a ticket's logical movement through a series of repeated steps. At the beginning of an investigation, a ticket can spawn in the queue of the Identification team, and the assigned investigator can review the indicators or evidence that initiated it. They may assign a task to collect data from a suspect machine. Or, after searching through past tickets, they may immediately conclude that a breach has occurred and send the ticket to the Containment team's queue. This team can receive the ticket and begin acting to restrict the attacker's movements.

The ticket can progress thusly through the workflow, from team to team and queue to queue, until it has been completed. Using workflows ensures that the right steps are performed in the right order, by the right staff.

2.2.2. Track all important data

Incident investigation involves collecting, interpreting, analyzing, and acting upon complex data in many forms. This data can describe the attacker, the target user, a system, data, an application, programs, code, communication methods, networks, as well as many other aspects of the computing environment or business operations. The data that is stored in the tracking system should not remain solely at the discretion of the

investigator. One investigator may believe certain data is more important to record than another, and as a result, material critical to a successful investigation may be lost.

The template for a master incident ticket should be populated at each step with the fields that the CSIRT, and the organization as a whole, see as important. Some fields can be marked as required, so the ticket cannot progress to the next step without the recording of certain data. Examples of useful data include:

- Assets that are compromised as well as suspected of compromise, their geographic location, owner, operating system, hardware model and details, and the indicator of compromise that led to their inclusion;
- Data that was stolen or suspected of being stolen, its storage location, owner, and classification level;
- Information about the threat actor, including the accounts or tools they are using, remote hosts they are operating from, and potential motive;
- Results of analysis, from reverse engineering of malware or tools, forensic evidence and artifacts discovered;
- Links to internal knowledge bases, or third party information sources about the attacker, tools, vulnerabilities, or stolen data.

Access to some fields may need to be restricted, so that their inadvertent dissemination does not expose the organization to further harm. An example would be the details of data stolen by an attacker like a social security number, credit card number, or credentials.

Some data may need to be set up for easy export, so that it can be used for detection/hunting or remediation. Examples might include hashes of malware samples, registry keys or other endpoint artifacts, or IP addresses used by attackers for command and control.

Justin Hall, justin.hall@cbts.net

2.2.3. Make searching easy

Few issues slow an investigation's progress as much as searching. Investigators researching a lead or trying to answer a question using notes on paper, or even individual document files on a computer, could spend minutes or even hours looking for one key piece of information. Modern search algorithms can be employed to render this problem irrelevant.

Most incident tracking tools can be configured to index some or all of the fields in an incident ticket. The goal to which a CSIRT should aspire is for an investigator to be able to search for any string of text in any field of all tickets. An investigator may know a domain name, for example, and want to know if it has been used in any previously known attacks. This data may exist in a field specifically created to store command and control hosts, but it also may exist in a snippet of a log file stored in a general 'notes' text field. Saving an investigator's time by allowing a full text search will greatly improve efficiency.

Allowing users to 'tag' tickets with specific keywords may also improve search capability. If a threat actor is referred to colloquially using a code name by the CSIRT or the security field at large, the code name could be used as a tag for an incident ticket in which the actor was suspected to be involved. Or, if several incidents involve the theft of schematic drawings from an organization, the tag 'schematics' could be applied to all relevant tickets. Providing the investigators with an automatically generated list of tags from which to search makes this feature even more useful.

2.2.4. Regular review of incident tickets

While an organization's CSIRT members may be exceptionally skilled at forensics and incident response, it is possible they may not all be well organized or motivated to document completely every time. As a result, CSIRT leads may want to

regularly review the content of an incident ticket to ensure investigators are following processes, avoiding shortcuts, or doing lazy work.

For example, a repetitive incident investigation that happens several times a week may become routine or boring to an investigator. The investigator may proceed through the ticket workflow, but after several cases, begin to fill in mandatory fields with incomplete data simply to bypass the documentation requirement and move the ticket along. In this situation, a CSIRT lead will need to correct the investigator's behavior, to ensure that all incident tickets have as complete of data as possible, every time. CSIRTs will quickly learn that today's useless data often becomes tomorrow's critical data.

This also may present an opportunity to revisit policies that dictate which data is mandatory and lighten the burden on the investigators.

2.2.5. Regular training

Complex incident response processes and complex investigations lend to a complex incident tracking tool. New investigators may not totally understand the tool's features or the CSIRT's implementation of their IR processes and procedures in the tool. Regular training for the team on the correct usage of the tool should be conducted, especially as new staff joins the team or as new features or changes are implemented in the tool's functionality.

This also gives team members an opportunity to ask questions that may have been present in their minds for some time and a chance to suggest new features or changes that have been necessary.

2.2.5. Operational security

A smart intruder will know their target, and many will immediately locate the tools and knowledge bases used by the information security teams employed by the target. Incident tracking tools will, by definition, be a key area of interest for such an

Justin Hall, justin.hall@cbts.net

attacker. They may want to simply know what you know about them; they may want to disrupt your ability to respond, record data, or modify the data stored to mislead the investigation.

Operational security measures should be taken to protect the confidentiality, integrity, and availability of the tracking tool. Data should be backed up daily and include the operating system as well as the tracking application and its database. Backups should be encrypted, and recovery capability should be tested regularly. (SANS CSC 8, 2015)

Incident tracking tools are usually built from software and as such require patching. Updates should be applied as the vendor releases them. (SANS CSC 4, 2015)

Access to the tracking tool should require strong authentication for users and administrators. In a client/server architecture, sessions with the server application should be encrypted. (SANS CSC 17, 2015)

It may be necessary to operate the incident tracking tool out-of-band in cases where attacker interference is suspected or expected. The tool could be operated on a separate physical network, to which a different physical client is connected or with which a VPN connection must be used for access.

2.3. Incident tracking tools

Finding a tool to use for incident tracking can be challenging. A CSIRT may not have budgetary resources, or the skillset to deploy or maintain a tool. A thorough review of the team's requirements should precede any attempt to acquire, design, deploy, or use a tracking tool.

2.3.1. Collaboration software

Many organizations resort to spreadsheets, such as Microsoft Excel, or other widely available standalone productivity software when doing ad-hoc organization or

tracking of activity. While this is a common practice, it may not be an effective means of maintaining records of incident investigation. They are typically too complex, with too much data to record, and do not lend well to collaboration due to the asynchronous nature of files that can only be opened by one user at a time.

Wiki software (such as MediaWiki or Atlassian Confluence), which allows ad-hoc content management by multiple simultaneous users, can often be a good starting place for CSIRTs looking for a low-maintenance, easy-to-deploy incident tracking system. Wikis tend by default to meet collaboration and search requirements, but most do not offer any kind of workflow capability. That said, wiki software is typically simple to deploy in an on-premises or hosted solution, and does not require much administrative overhead to keep operational. User accounts can be added and then investigators may begin to create wiki pages for each incident, using free-form documentation or template pages that encourage investigators to document specific information.

2.3.2. Support and issue-tracking software

By default, many organizations track incident investigations using help desk or support tickets. This practice most often happens when a user-reported issue, such as a performance problem, develops into a security issue after a period of troubleshooting. The 'ticket' framework present in help desk and service software, which allows multiple technicians to update a single ticket, track its state, and move it through a workflow, may make this kind of tool (such as Remedy or ServiceNow) attractive for incident tracking purposes.

However, the inability for many help desk solutions to allow customization of fields, workflows, and searching could impede effectiveness at incident tracking. Exporting specific data for use in security monitoring systems may be challenging, as it is

not typically a use case that is necessary when using a help desk solution for its intended purpose.

Bug or issue tracking systems, used by software development teams, are similar in concept to help desk software, in that both attempt to use individual cases or tickets to track effort and activity in the solving of certain technical problems. However, the inherent need developers have for custom workflows, issue types, and fields, makes an issue tracking system, such as Bugzilla or Atlassian JIRA, a better fit as an incident tracking tool.

A CSIRT inside an organization with a software development team could use that team's issue tracking product and set up a custom project or instance for their purposes. This could allow the CSIRT to deploy an IR workflow, users, queues, and ticket templates with custom fields that exist separately from those used by the developers.

2.3.3. Purpose-build Incident Tracking Software

Of course, attempting to shoehorn incident tracking capability into a product for which it was not designed could be a time consuming and frustrating effort for a CSIRT, especially when software designed solely for this use is available. CSIRTs should start with a purpose-built tool, if resources allow, to minimize start up time and customization work. Unfortunately, there are few purpose-built tools available.

Best Practical's RTIR is a free, open-source, web-based incident tracking tool used by several CSIRTs. The tool began as Request Tracker, a bug and issue tracking system, and the incident response features were later added as a separate module. While an open-source product, technical support is available from the project's maintaner, Best Practical. ("RTIR: RT for Incident Response - Best Practical," n.d.). The low cost to entry makes RTIR a good starting point for CSIRTs that are deploying their first incident tracking system, but are operating on limited budgets. The tradeoff comes from the IT

operations resources necessary to install, customize and manage the appliaction and its components, which could be substantial.

AIRT, or Application for Incident Reponse Teams, is another web-based incident tracking tool ("Application for Incident Response Teams (AIRT)," n.d.). AIRT is relatively simple in features and functionality compared to RTIR; however for a small team with few incidents to track per year, it may suffice. It does not appear to be maintained, though: the software has not been updated since 2009, and links for documentation lead to web servers that are unavailable.

A commercial alternative is CyberSponse's Security Operations Platform. This product claims to offer all of the features necessary to perform effective incident tracking and can be used relatively quickly by a CSIRT after initial deployment. Customers can deploy the tool on-premises or in a hosted environment ("IR Team," n.d.). While the tool is purported to be feature-rich, it is relatively new and untested as of this paper's publishing, and if its cost is substantial, it may be a deterrent to resource-limited CSIRTs.

3. Conclusion

Even the most mature CSIRT can function without a formal incident tracking system, but it will find itself struggling with managing tasks, effective collaboration, and repeating work between different incidents. The benefits of using a tracking tool far outweigh the difficulties.

Standing up an incident tracking tool is much like the deployment of any other server application. While requiring ongoing maintenance and administration, if configured properly and designed to mirror the team's IR processes, the tool can be a substantial help to investigators, IT operations teams, CSIRT leads and risk managers.

Understanding and documenting the CSIRT's IR process at the outset is key. Configuring workflows, queues, and custom fields according to the steps that must be

followed and the data that must be collected should happen before the tool is used. Regular updates to the tool should be based on feedback from its users and growth in experience and knowledge of the team.

The system itself could be a wiki, bug tracking tool or a purpose-built incident management product. It must serve the mission of the security group, while not exhausting the resources of the organization, CSIRT, or IT support. It must be protected from misuse and defended against compromise by an attacker. If this can be accomplished, it is certain that the instutition of formal incident tracking would be considered a positive achievement in the growth of a CSIRT and will contribute to the overall security posture of its organization.

Justin Hall, justin.hall@cbts.net

References

Cichonski, P., Millar, T., Grance, T., & Scarfone, K.. (2012). *Computer security incident handling guide: Recommendations of the National Institute of Standards and Technology*. Gaithersburg, MD: U.S. Dept. of Commerce, National Institute of Standards and Technology.

Gibson, D. (2011). *Managing risk in information systems*. Sudbury, MA: Jones & Bartlett Learning.

Skoudis, E. (2010). Incident handling step-by-step and computer crime investigation. *SEC504: Hacker tools, techniques, exploits and incident handling* (p. 14). Bethesda, MD: SANS Institute.

SANS Institute - Critical Security Control: 8. (n.d.). Retrieved from http://www.sans.org/critical-security-controls/control/8

SANS Institute - Critical Security Control: 4. (n.d.). Retrieved from http://www.sans.org/critical-security-controls/control/4

SANS Institute - Critical Security Control: 17. (n.d.). Retrieved from http://www.sans.org/critical-security-controls/control/17

RTIR: RT for Incident Response - Best Practical. (n.d.). Retrieved from https://www.bestpractical.com/rtir/

Application for Incident Response Teams (AIRT). (n.d.). Retrieved from http://airt.leune.com/

IR Team. (n.d.). Retrieved from http://cybersponse.com/ir-team

Justin Hall, justin.hall@cbts.net