south dakota bit bureau of information & telecommunications

CYBER LUMBERJACKS & LOG CONSOLIDATION

SouthDakota

CYBERSECURITY
STATE OF SOUTH DAKOTA
JIM.EDMAN@STATE.SD.US

FEBRUARY 2018 - DECEMBER 2018

Improving Endpoint Visibility for Enhanced Incident Response at the Speed of Search

EXEMPLAR

The everyday goal of all cybersecurity teams is to ensure that their organization's confidential digital assets are secured. The mention of the terms breach or data leak ruin the day of any cybersecurity professional. The attack vectors used by bad actors continue to grow. Application software is a popular target, social engineering and the human firewall will always be a high risk, supply chain security has become a great concern recently, and finally endpoint protection has been a primary target since the first computer virus was written in 1986. Whether the endpoints are desktop computers, laptops, servers (physical and virtual), smartphones, printers or other handheld devices, protecting and analyzing the endpoint devices has long been a challenge within the industry.

The technology aspect of endpoint protection has improved drastically over time. Beginning with the definition-based signatures of anti-virus products, solutions today offer much more features. In addition to recognizing common file-based exploits, desktop protection now includes active scanning services, behavior-based protection, and URL sandbox detonation capabilities.

Incidents must be contained as quickly as possible to limit their impact. Complete recovery from an incident requires knowing the original attack vector and how to bolster the defenses for the specific point of compromise. The defensive solution must then be applied on an enterprise-level to avoid a repeat compromise.

Incident Response includes the tenants of Prepare; Identify; Contain; Eradicate; Recover; and Review. The common piece missing from our solution set is a cohesive solution to forensically piece-together what happened once an infection occurs. This project targeted aspects of the Contain and Recover phases of Incident Response, answering such as:

Where was the point of entry for the compromise / infection?
What vulnerability was exploited?
What is the scope of the compromise?
Was any data exfiltrated?

During normal monitoring tasks and suspicious events, analysts must be able to address the situations rapidly at a scale of thousands of workstations while under great pressure, knowing the potential consequences at stake. From this comes the need for greater endpoint visibility and improved forensic analysis. As a highly centralized information technology organization, the Bureau of Information and Telecommunications (BIT) serving South Dakota state government supports more than 9,000 Windows endpoints that span across 36 agencies. While admittedly small compared to other governments, our size also creates the opportunity to implement a creative solution on an enterprise level.

Being able to review any nefarious activity on a workstation in real-time or playback mode is a transformational stride forward in the security realm. Adding the workstation analysis aspect to current network-based monitoring capabilities is a revolutionary step forward. Analysis of network and application forensic data has long been a common service done by security analysts. Taking that ability down to the workstation across an enterprise with real-time analysis capabilities is a monumental leap forward in security services!

bit

CONCEPT

A history of inefficient incident response and partial forensic investigations for compromised systems was the genesis for the project. The South Dakota BIT Security Operations Center does an excellent job identifying threats as alarmed by security alerts and malicious emails. Unfortunately, the response to the events was slower than desired because of the lack of cohesive analysis. Some technologies had no residual data gathered, while other technologies stored different aspects of the event. This gap prevented a complete view of any events.

Investigations involve a variety of technology including software applications, network appliances, and logs from endpoints. Investigating security alerts at the network level are not conclusive to understand the endpoint threat so additional investigation was needed. Obtaining logs from the endpoint is a tedious process that involves remotely connecting to the system, finding the relevant logs, exporting the logs, and finally reviewing the logs. This manual, intensive process can yield enough useful information to determine if the workstation is operating normally or if it is compromised. The price paid for this lengthy process is time, though. The baseline performance of a workstation varies drastically dependent on the work requirements of the individual, negating a comparison to any performance standards. If a workstation's logs prove that the system is compromised, there is no way to identify other workstations with the same logs without repeating the manual intensive process by connecting to each machine. Thus, the need to consolidate every relevant endpoint log was clear. A system needed to be built to consolidate security event logs for every computer, and it had to be searchable.

Scenario: The Security Operations Center was alerted by the enterprise network firewall that a Windows 10 workstation was trying to establish an outbound connection to a known botnet IP address originating in Jamaica. The SOC analysts initially reviewed the network artifacts that the endpoint was generating. All appeared to be normal. The analysts then acquired the victim desktop's Windows logs and began searching around the timeframe that the firewall detected the traffic. The Windows logs revealed that user's last actions on the system before the connection was that Outlook created a process called Word.exe which then opened powershell.exe with the command line arguments that contained the IP address noted from the firewall log. These Windows logs provided enough information to understand that the user opened an email that contained a malicious Office document that tried to download malware from Jamaica.

While the analysts above were able to find the source of the threat to
that device, there were three assumptions:

1. An alert had to be generated to identify the threat;
2. All timestamps across monitoring services must be synchronized across multiple technologies;
3. The Windows logs must be accessible, requiring physical or virtual access to the machine.

It was clear that there was a visibility gap for SOC analysts that needed to be filled.

The following questions were driving factors.

What if all the endpoint logs could be centrally gathered in real-time for searching and analysis?
What if the SOC did not have a firewall or network alert to start an investigation?

The gears started turning for our security engineers to set forth and start the endpoint visibility project. The team thought if the endpoint logs could be stored alongside other security events presently gathered in their Security Information and Event Management (SIEM) platform, they could add a significant missing piece to our security posture. However, upon a proof of concept and vendor quotes, it would have costed upwards of $200K in software, hardware and support that was not budgeted. While many products exist in the Endpoint Detection and Response space, there were far too many that were fiscally out of reach. The State of South Dakota settled on using open-source technology that was fast, scalable, and easy to use.

The goal became to build a 'single pane of glass' for viewing all endpoint security logs and establish more efficient threat analysis, incident response, and forensics. A 'DVR-like' system would allow security analysts to walk through security incidents with enough information to determine the root cause of a compromise and identify any other systems affected across the enterprise network.

The architecture requires three primary parts.

1. The first part is the event log shipper that resides on the workstation, which in our case runs Windows 10.

2. The second part is the database that will store the event logs.

3. The third part is the front-end user interface to search the aggregated logs in a useful manner.

This interface is how analysts interact with the terabytes worth of data.

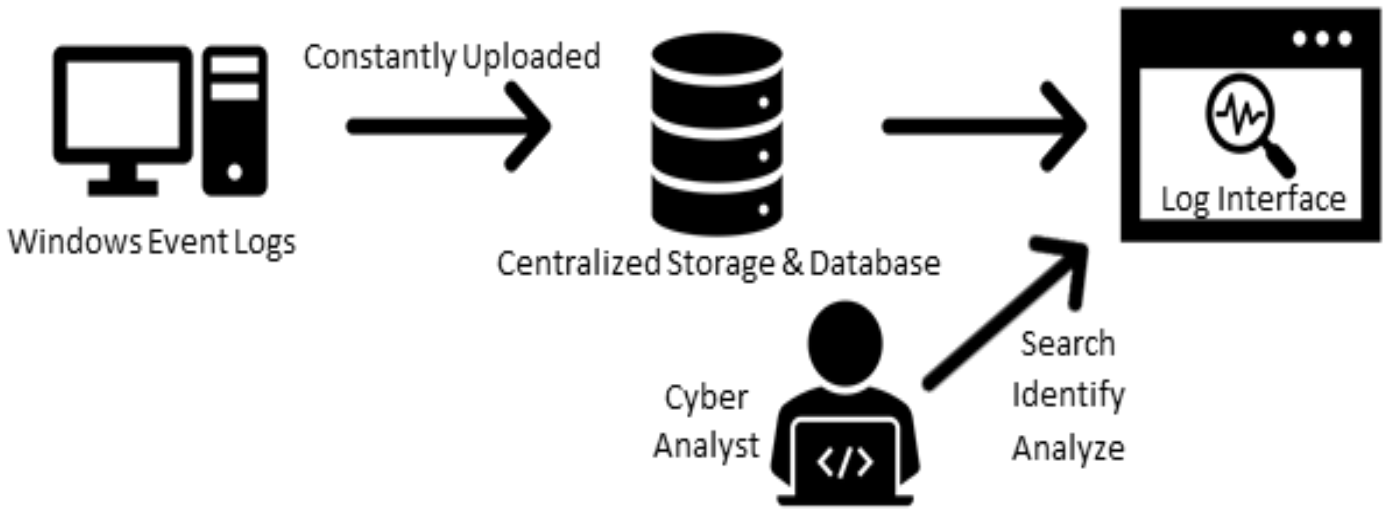

SIGNIFICANCE

There are five key items to the significance of this project for the State of South Dakota:

1. Improved alerting- In the past, our response to security incidents relied on anti-virus definitions, firewall recording of events, and intrusion detection/prevention systems. Adding endpoint visibility to this list provides the detailed data from the point of compromise.

2. Faster incident response times- The SOC moved from a manual, intensive log gathering process to a central, searchable repository for thousands of workstations. Incredibly, a search across every workstation for a 30-day time window takes only a few seconds. Previously, this type of information at the analyst's fingertips would take at least an hour for a single machine, and days for a wider group of machines.

3. Improved after-action reports- The granularity of log data helps analysts build the timeline of the events that occurred on a host system from before the malicious activity happened up to the time the system was taken offline during the incident response process. Did the threat start from a malicious email? Did the end user download a malicious application from the Internet? These questions are easily answered with all the relevant details which make for great reports.

4. Rapid response to enterprise impact- Solving a single workstation issue is relatively simple. Delivering that solution on an enterprise level is much more complex. Being able to identify other workstations with the same behavior is easily discoverable. Regardless of whether that solution involves a software update or a new security layer, the endpoint logs provide the empirical evidence for the next recovery steps.

5. Expanded visibility- Today's mobile worker requires access to their digital resources from anywhere, including non-trusted networks. This creates a gap in our security posture, as we were unable to receive the security logs in real-time. However, once a system has returned to the enterprise network, the logs are forwarded to the consolidated log source for review. This aids analysts with what activity was performed while the system was connected to potentially dangerous public networks.

IMPACT

The obvious impact is the ability to quickly search over 9,000 computers for security events. It is an impossible task for a person to view the logs without becoming overwhelmed. Building an effective interface to the logs required a dashboard to allow for searching any fields such as computer name, user name, or even a process name. Searching for a specific malware process name must be easy and quick including the ability to identify any entries with specific file names. A dashboard with instantaneous results of interesting or actionable events is highly valuable for real time analysis. The dashboards must include pie charts, line graphs, and other visuals that provide insight about the consolidated logs.

bit

Process activity with network connections outside the US may be an indicator of malicious activity. This visualization allows analysts to see where computer processes are communicating.

T

Our people: the SOC can identify and remediate incidents quicker with more information.

Our clients: keeping technology services more dependable without interrupting their job functions.

Our constituents: improved protection of their data and continued uptime during cybersecurity incidents.

With any sort of investigation (cyber, criminal, etc.), the more details a detective has, the greater their ability to solve the incident. Having the endpoint log data in an easy to use system provides cyber analysts the data they need to investigate anomalous activity. This allows us to recreate the malicious activities without relying on the users, taking an objective analysis of what happened. Gathering and consolidating the endpoint logs in near real-time allows the SOC to act quickly when an incident occurs. After identifying malicious activity, a thorough incident response process begins. Having the context of the malicious activity is crucial to understand what the root causes are for the affected system and what other technology entities need to become involved.

Once the involved teams are notified of an incident, actions to mitigate the threat start by removing the device from the network. This is where the impact to the state agencies comes into play. Before this project, the employee was unable to use their computer until a full forensic analysis had been completed on the system. This forensics analysis would require the physical system to be accessed by the SOC to gather the required logs. The forensic process to manually gather these details took many hours, sometimes days due to physical location or network speeds of the device. While some remote forensics were possible, it was still very difficult to gather logs from systems that had a very limited network connection that wouldn't allow for large log extractions.

During these times of forensic log extraction, the employee's computer would be unavailable to them unless the agency was fortunate enough to have a spare (not likely). Eventually, all logs would have been harvested and the afflicted system gets its hard drive erased and the operating system and the employee's apps will be re-installed.

The forensic extraction of logs that took hours and sometimes days has been greatly reduced to minutes. The analyst's ability to quickly search the logs helps aid them in differentiating between false and true positives of threat signatures. Specifically, when the enterprise firewall systems report a device is exhibiting known malware traffic, a search across the endpoint logs provides confirmation or disaffirmation. This immediate feedback avoids the extensive forensic processes to discover that nothing malicious happened and the signaling event was only a false positive.

Dependable I/T services add great value in providing productive state employees. Cybersecurity in state government is a tenuous balance between the requirements of open government and securing our confidential information. Employees and constituents are migrating from the typical 8:00 – 5:00 workday to 24x7x365 availability. This requires that our service models minimize the impact of power outages, technology failures, fiber cuts, and cybersecurity incidents. Any incident that takes an employee offline, results in a negative service impact to our constituents. Those services impacted could include a driver's license request, vehicle registration renewal, criminal investigation or a benefits program not being delivered. The obvious initial goal is to avoid any cybersecurity event but after an incident occurs, the critical needs become Identify; Contain; Eradicate; Recover; and Review. Original products evaluated came with project costs exceeding $200,000. That was a figure not in our budget. Our open-source solution, including training and implementation, came in under $35,000, which was a considerable cost-avoidance. From a staff performance and efficiency perspective, the consolidated log processing and analysis project delivered a vastly improved incident response process and ultimately a more cyber-secure state government.