Analyzing Risks of Virtual Private Network Connections

David Daggett

Abstract

The use of Splunk for analyzing VPN logs is an effective approach for identifying vulnerabilities in network endpoints. Splunk, a powerful platform for searching, monitoring, and analyzing machine-generated data, enables organizations to aggregate VPN logs in real-time, providing insights into network activity, user behavior, and potential security risks. By indexing VPN traffic and authentication logs, security teams can track abnormal patterns such as multiple failed login attempts, unusual IP addresses, or unexpected changes in bandwidth usage, all of which could indicate potential vulnerabilities or breaches.

With Splunk's advanced search and reporting capabilities, users can create custom dashboards and alerts to detect suspicious activities. Automated searches can flag endpoints exhibiting unusual behavior, while correlation analysis can identify links between compromised devices and broader network vulnerabilities. In particular, Splunk's machine learning capabilities can be leveraged to predict and prevent threats by identifying trends that might otherwise be missed in traditional log analysis. This proactive approach to monitoring VPN logs allows for the early detection of security weaknesses, enabling rapid response and minimizing potential damage to network integrity. By enhancing endpoint visibility, Splunk plays a crucial role in securing remote connections and safeguarding sensitive information.

Additionally, Splunk's automation and alerting features allow teams to create custom workflows that notify them of vulnerable or misconfigured endpoints identified through Shodan. This synergy between Splunk's log analysis and Shodan's device intelligence enhances an organization's ability to proactively identify and mitigate security risks, improving the overall resilience of their VPN infrastructure.

Introduction

A Virtual Private Network (VPN) creates a point-to-point tunnel that encrypts your personal data, masks your IP address, and lets you access internal parts of a network that would normally be blocked by a firewall [1]. The traffic traversing across a network at any given moment can range in the tens of thousands of packets. This high volume makes it impossible to identify malicious traffic without protocols in place, but with the use of tools like security information and event management (SIEM), this large amount of data becomes manageable. Splunk is one of those tools that helps organizations detect, analyze, and respond to security threats from data collected from different log sources. The focus of this project was to investigate devices connecting via VPN and the country of origin of those connections then create a visual dashboard of the information.

Progress

When I started my project there were many alerts with a focus on VPN traffic, but none looked at the vulnerabilities from endpoints or connections from sensitive countries. A "sensitive country" is a country that receives special consideration for policy reasons, such as national security, nuclear nonproliferation, regional instability, threat to national economic security, or terrorism concerns. Using the Splunk indexes already available, I filtered out all failed connections to the network and any IP addresses that were internal.

FirstSeen	LastSeen	ExternalIP	Count	User	Port	Vulnerability
11/11/2024 3:24	11/12/2024 12:37	155.209.124.240	177	user1	668	CVE-2023-44487
11/10/2024 17:46	11/13/2024 15:43	249.102.36.186	188	user2	4803	None
11/9/2024 16:57	11/12/2024 22:53	151.162.102.142	169	user3	8930	None
11/9/2024 18:04	11/13/2024 4:25	225.15.20.188	72	user4	7346	None
11/9/2024 23:47	11/12/2024 19:30	226.219.17.59	219	user5	2744	None
11/9/2024 5:14	11/13/2024 0:26	115.127.167.223	21	user6	8839	None
11/9/2024 0:32	11/12/2024 20:48	60.201.89.143	204	user7	1898	None
11/11/2024 11:23	11/12/2024 14:38	19.244.215.232	249	user8	8983	CVE-2018-11759
11/12/2024 23:08	11/13/2024 5:05	76.180.122.142	167	user9	7383	None
11/10/2024 10:37	11/13/2024 1:11	124.48.115.7	29	user10	4264	None
11/9/2024 10:28	11/12/2024 17:16	99.150.47.129	44	user11	2827	None
11/10/2024 11:20	11/13/2024 12:56	86.110.66.120	131	user12	8431	None
11/8/2024 17:25	11/13/2024 5:13	242.76.96.10	36	user13	5866	None
11/9/2024 0:22	11/11/2024 4:27	61.37.227.5	60	user14	2879	None
11/8/2024 21:56	11/11/2024 0:02	29.217.40.238	204	user15	6417	None
11/9/2024 7:09	11/13/2024 2:54	62.216.216.5	90	user16	9334	CVE-2023-44487
11/9/2024 9:35	11/10/2024 13:35	162.247.6.112	199	user17	2142	None
11/11/2024 8:52	11/12/2024 4:42	240.104.209.61	190	user18	8879	None
11/8/2024 23:47	11/10/2024 6:07	220.12.251.97	222	user19	6035	None
11/9/2024 21:59	11/12/2024 1:03	113.45.214.57	223	user20	7083	None

Figure 1 – Example VPN traffic logs. The external IP addresses within the image are mock IP addresses and were not captured from current PNNL staff/traffic.

The information is returned in packets containing most of the data in a single field that needs to be split up to be more easily read. By using a series of regex commands, we were able to pull out useful information such as who the users were, the external IP they were connecting from, and what internal connection they were using. The connections produced that were infrequent were less of a concern as it could be attributed to a connection from a temporary location. We could

reduce the number of results to only show connections that had a higher frequency of appearance based on this conclusion. To get a better overview of the VPN network, we expanded the search to take the connections with a high count in a 24-hour period then search for those same connections over a week time frame. These findings presented us with solid data set to evaluate. The IP addresses found from the search were run through Shodan to find what ports were open on these endpoints and what known vulnerabilities were associated with them. This data presents opportunities for mitigation before an incident and could prevent system compromise. Additionally, we wanted to investigate whether there were connections from sensitive countries. Using the same index as the prior search, we could filter the results using a predetermined list of countries. After gathering information, we created two different dashboards to visualize and highlight key parts of the data.

Future Work

Currently, the speed at which Shodan processes inputs is quite slow. Optimizing either the queries used, or how Shodan operates, would improve the speed at which we can check each day for new inputs. There are some additional panels that would benefit being added to the dashboards, such as creating a panel that focuses on identifying brute force attacks and which users have been approved for travel and their destination countries. The information from these additional panels could provide insight into whether the connection originated from an actual user.

Currently, any vulnerability found is returned as a common vulnerability and exploit (CVE) reference number. The dashboard does not provide any additional context for the CVE. In the future, a Splunk lookup can be leveraged to have two columns; one that displays the CVE, and one that provides a brief description of the CVE. This will eliminate the lack of context and give quick insight into what vulnerabilities exist.

Impact on Laboratory or National Missions

PNNL's mission is to transform the world through courageous discovery and innovation. Fully embracing that mission, researchers collaborate with local, state, and federal government agencies, academia, and industry partners. Theses collaboration efforts may produce data that may be considered sensitive and cannot be released to the public. The additional alerts and capabilities mentioned here can help to build upon existing information system continuous monitoring which will help to continue protecting data and assets. The alerting enhancements here also helps to quickly identify any potential points of compromise and to implement timely actions to address before an incident. These actions help to continue to support the mission statement of PNNL and increase the security posture of the ongoing continuous monitoring efforts.

Conclusions

The focus of this research was to analyze logs related to VPN connects and the endpoints associated with a public IP address. After reviewing the information returned from the alerts, a minimal number of vulnerabilities were seen from endpoints. The continued use of these alerts may not be needed but shows no drawback in continued use. VPNs take a considerable amount of effort for adversaries to bypass or compromise. The various layers of defense-in-depth cyber security enhance the security profile of PNNL's VPN, regardless of the location from which users connect.

References

[1] Bansode, R., & Girdhar, A. (2021, January 1). *IOPscience*. Journal of Physics: Conference Series. https://iopscience.iop.org/article/10.1088/1742-6596/1714/1/012045

[2] S. Nandhini, Shubham Vaishnav, S. Vikas Vishwakarma; VPN blocker and recognizing the pattern of IP address. *AIP Conf. Proc.* 5 April 2022; 2405 (1): 030019.

Appendix

Name	Institution	Role
David Daggett	Pacific Northwest National Laboratory	Primary.
Jacob Kromm	Pacific Northwest National Laboratory	Mentor. Provided feedback and troubleshooting.
Nancy Roe	Pacific Northwest National Laboratory	Acted as support for CCI Internship activities and deliverables.
Aaron Haglund	Pacific Northwest National Laboratory	Document review.

Acknowledgements

This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists (WDTS) under the Community College Internship Program (CCI).

Notable Outcomes

Research will be presented at PNNL's December Virtual Research Symposium on December 12, 2024.