Field Note on CVE-2019-11510: Pulse Connect Secure SSL-VPN in the Netherlands

MATTHIJS KOOT, Secura BV & University of Amsterdam

This Field Note describes the case of a critical unauthenticated RCE vulnerability in an SSL-VPN product that remained unpatched at a large scale-up and until after exploits became public. Approximately 14,500 systems worldwide were reportedly unpatched at the end of August 2019. Two weeks after exploits emerged in public, both GCHQ and NSA released notices that the vulnerability was being exploited by APT actors. The present Field Note describes observations from the Netherlands and includes reflections in an attempt to stimulate thinking on how to improve the status quo, such as through coordinated proactive measures by CSIRTs.

CCS Concepts: • **Security and privacy** → *Vulnerability scanners*; **Network security**; *Intrusion detection systems*;

Additional Key Words and Phrases: Situational awareness, large-scale Internet scanning, vulnerability detection, threat intelligence

ACM Reference format:

Matthijs Koot. 2020. Field Note on CVE-2019-11510: Pulse Connect Secure SSL-VPN in the Netherlands. *Digit. Threat.: Res. Pract.* 1, 2, Article 13 (May 2020), 7 pages.
<https://doi.org/10.1145/3382765>

1 INTRODUCTION

In March 2019, Orange Tsai and Meh Chang¹ discovered critical vulnerabilities in Pulse Connect Secure SSL-VPN.² The researchers reported the vulnerabilities to Pulse Secure on March 22, 2019 [11]. The vulnerability that was assigned CVE-2019-11510 (CVSSv3 Base Score: 10) constitutes an unauthenticated remote path traversal that allows attackers to obtain credentials (cached usernames and passwords, including users who are authenticated via an LDAP interface, e.g., Active Directory) and active VPN session identifiers. The latter can be used to hijack active multifactor-authenticated (MFA) VPN sessions, hence bypassing MFA.

On April 24, Pulse Secure released security updates (SA44101) and urged customers to apply those [8]. In early August, the researchers detailed their findings at hacker conferences Black Hat USA (August 7) and DEF CON (August 9) but did not release a proof-of-concept (PoC) exploit at that time.

On August 21, a hit/no-hit PoC for CVE-2019-11510 was published.³ On August 22, CVE-2019-11510 was reported to be seen being exploited in the wild [1]. Two weeks after exploits emerged in public, both GCHQ and NSA released notices that the vulnerability was being exploited by APT actors [5, 6].

¹ Researchers from the Taiwan-based organization DEVCORE: <https://devco.re/en/about/>.

² In an effort that started in December 2018, the researchers discovered critical vulnerabilities in multiple SSL-VPN products: FortiGate, Palo Alto GlobalProtect, and Pulse Connect Secure. This Field Note focuses on Pulse Connect Secure.

³ Website: <https://www.exploit-db.com/exploits/47297>.

Authors' Address: M. Koot, Secura BV, attn. M. Koot, Vestdijk 59, 5611 CA, Eindhoven; email: koot@uva.nl.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2020 Copyright held by the owner/author(s).

2576-5337/2020/05-ART13

<https://doi.org/10.1145/3382765>

Fig. 1. Systems vulnerable to CVE-2019-11510 in Dutch-registered IPv4 space over time.

An assessment of Dutch-registered IPv4 space, carried out on the weekend of August 24⁴ by the author, yielded 538 systems vulnerable to CVE-2019-11510. The systems were traced to some 200 organizations via WHOIS data from LIRs and RIRs, DNS records, hostnames present in the TLS certificates, customization of the login page (organization name/logo), and information collected about surrounding IP addresses. The scope of affected organizations included Dutch local governments and a national government department, defense industry, financial sector, aerospace sector, education, healthcare, IT providers, petrochemical industry, and harbor and transport organizations, among others. In many cases, the systems were production systems.

The plot in Figure 1 depicts the number of vulnerable systems on Dutch-registered IPv4 space over time. Initial target IPs and IP ranges were obtained from commercial sources including Shodan and BinaryEdge, public sources such as RIPE NCC, and private/undisclosed sources. These were probed using ZMap, Nmap, and custom scripts. All IPs running Pulse Connect Secure were then tested for CVE-2019-11510 on tcp/443 by exploiting the path traversal to read the file “/etc/passwd”⁵ and then searching the output for the string “root:x:0:0:root,” the presence of which indicates vulnerability: [https://\[TARGET\]/dana-na/..../dana/html5acc/guacamole/..../..../etc/passwd?/dana/html5acc/guacamole/](https://[TARGET]/dana-na/..../dana/html5acc/guacamole/..../..../etc/passwd?/dana/html5acc/guacamole/).

For monitoring purposes, all systems that were identified as vulnerable during the initial scan were tested again on a daily basis. The number of vulnerable systems decreased by approximately 60% in 10 calendar days following efforts by multiple parties. For a vulnerability as critical as CVE-2019-11510, a much shorter decrease slope is desirable, especially once a PoC is publicly available or anticipated to become available soon. Organizations that remained vulnerable to CVE-2019-11510 included many that are attractive to attackers.⁶ Table 1 lists the events over time.

⁴Later that same day, Bad Packets published the results of a global-scale test that yielded 14,500 vulnerable systems worldwide [7].

⁵On Pulse Connect Secure systems, that file does not contain any user data: it is identical across different systems that run the same version of Pulse Connect Secure. Hence, no user data is compromised by the test. Nonetheless, the test in principle constitutes a violation of Dutch computer crime legislation.

⁶Vulnerable instances of Pulse Connect Secure are implicated in the ransomware infection that affected Travelex in late 2019 [10]. Travelex has stated that the patch was installed in November, more than 2 months after a PoC exploit became public.

Table 1. Timeline of Events

Date	Event
2019-03-08	Initial discovery of vulnerabilities
2019-03-22	Vulnerabilities reported to vendor
2019-04-24	Vendor releases patches
2019-05-01	NCSC-NL Advisory NCSC-2019-0353 ⁷ : Medium × High
2019-08-07	Findings presented at Black Hat USA
2019-08-09	Findings presented at DEF CON 27
2019-08-20	Hit/no-hit PoC for CVE-2019-11510 released
2019-08-21	NCSC-NL Advisory NCSC-2019-0353 updated: High x High
2019-08-22	CVE-2019-11510 seen exploited in the wild
2019-08-24	Dutch IPv4 space scanned: 538 vulnerable systems
2019-08-25	Results shared w/NCSC-NL
2019-09-01	320 vulnerable systems remaining
2019-09-08	202 vulnerable systems remaining
2019-09-15	163 vulnerable systems remaining
2019-09-22	153 vulnerable systems remaining
2019-09-29	112 vulnerable systems remaining
2019-10-06	85 vulnerable systems remaining

2 REFLECTIONS

We take the case of the Pulse Connect Secure vulnerability to raise questions around societal vulnerability stemming from such systems. One such question is how deferred patching behavior can be explained.

In August 2018, a report by tCell stated that it took organizations “an average of 34 days” to patch critical vulnerabilities in production systems, which was “only 4 days faster than the average time to patch overall regardless of severity” [9]. As is clear from the timeline in Table 1, the patching of many Pulse Connect Secure systems took much longer—and in fact did not happen until the exploit became public. The organizations responsible for these systems did not install a critical patch in an Internet-facing system in 4 months. In a few cases, an explanation is known to the author, but a structured evaluation is much needed. This remains an open question for now.

Another question is how societies should cope with situations like this: critical vulnerabilities that may remain present in Internet-facing systems for longer periods, currently without anyone being aware of the scope of the problem. Individual organizations are themselves responsible for protecting their assets, but the present findings reconfirm the existing belief that additional efforts may need to be considered.

The root cause of the problem remains unknown. Explanations can include lack of situational awareness of one’s own Internet-facing infrastructure, inability to schedule downtime (production VPN servers may be considered “always-on” systems from a business perspective), difficulties in the patching process itself, and desensitization to “high × high” advisories (many such advisories exist, and not all are equal). The latter could be addressed through improvements in vulnerability/risk scoring methods that CSIRTs use.

A complementary possibility is to organize proactive capabilities—in addition to existing reactive capabilities—at the level of the individual CSIRTs in a federated model. In the Netherlands, multiple CSIRTs exist and new

⁷Source: <https://advisories.ncsc.nl/advisory?id=NCSC-2019-0353>.

CSIRTs are foreseen: besides the Dutch National Cyber Security Center (NCSC-NL) that focuses on the national government and vital infrastructure, a range of sector- and organization-specific CSIRTs exist for, among others, the financial sector, defense, managed service providers, municipalities, and healthcare. These CSIRTs are relatively close to their constituents and as such have a trust position that could be leveraged to provide confidential, targeted early warning based on active scans performed by the CSIRT on constituents' infrastructure—in other words, early warning based not only on threats but also vulnerabilities.

For parts of society that do not have a CSIRT, a CSIRT entity could be set up at the level of autonomous systems (AS) and be tasked along the same lines. After all, IP2ASN mapping is reliable, and abuse contacts are already known at the AS level.

In the EU, the NIS Directive⁸ prescribes that Member States should have a national authority that monitors the security of network and information systems, covering at least the sectors in Annex II and the services referred to in Annex III. It allows Member States to include other sectors and new types of measures. Proactive scanning could be such a measure.

In addition, as a separate topic, scanning data could be collected at the level of CSIRTs and be made available for independent scientific research. Data should be stripped from information that could lead to identification of organizations, and be released possibly under NDA and perhaps subject to mandatory security screenings. At the international level, one venue for data exchange between vetted parties is the IMPACT Cyber Trust⁹ initiative by the U.S. Department of Homeland Security. Sharing data about vulnerabilities—even after they have been fixed—remains a sensitive topic and should be approached with due care. Nevertheless, it would help “open the oyster” of information security and help shed light on a significant blind spot that digital societies face: that of an often unknown nature and scope of (critical) vulnerabilities across sectors. Longitudinal studies could then be performed, yielding insight that can help craft sector-specific measures to promote best practices and awareness. This is a way forward from the fragmented and sometimes anecdotal stories that are known to some but not documented or recorded in a way that allows cross-sector quantitative research.

Proactive scanning capabilities need not extend to continuous scanning of any and all systems and vulnerabilities. That would be intractable for CSIRTs and could disrupt organizations' self-monitoring activities, as well as the private information security sector. It suffices to organize, along the lines of existing CSIRT organizations and in conjunction with national implementations of the NIS Directive, a uniform way of testing for specific types of vulnerabilities, such as high-impact vulnerabilities that can be exploited unauthenticated remotely. Given the complex reality of CSIRTs and constituencies, the question whether or not this should be voluntary (opt-in/opt-out) for CSIRTs and constituents is to be carefully addressed—we leave it open for now.

For CSIRTs, proactive scanning means that they need a process to schedule time and priorities (in terms of possible impact, not all constituents are equal), and communicate with constituents. The latter can be time consuming and involve multiple exchanges between CSIRT and constituent, and sometimes vendors—and hence touch on aspects of relationship management. Constituents may desire the following:

- Transparency about the types of vulnerabilities for which their infrastructure will be tested by the CSIRT that covers them.
- The ability to easily distinguish network traffic of scans by CSIRTs from (possibly) malicious scans.
- The ability to opt-in or opt-out specific IP addresses or entire ranges of the constituent's IP block(s).
- A non-emergency contact at the CSIRT.

It may be a non-trivial effort, but we argue that it should be considered. Key questions that need to be addressed include legal, socio-ethical, and practical aspects of proactive scanning, such as follows:

⁸Website: <https://www.enisa.europa.eu/topics/nis-directive>.

⁹Website: <https://www.impactcybertrust.org>.

- Should (unsolicited) proactive scanning be performed at all?
 - Unsolicited scans have taken place for decades by good and bad actors alike and will continue regardless of our answer to that question. We plea in favor of proactive scanning under assumption of “no commerce, no press” conditions, as is usual in CSIRT realms.
 - Proactive scanning should not blur the constituent’s feeling of responsibility for its own digital infrastructure. By limiting the scanning to a no-guarantee early warning for critical vulnerabilities, the responsibilities on both ends remain clear and intact.
 - Proactive scanning should not introduce disproportionate risk to constituent infrastructure, such as the failure of system or network components. The CSIRT should take necessity, proportionality, and subsidiarity into account for every type of activity. The thoughts laid out in van der Ham and van Rijswijk-Deij [12] on ethics in the context of Internet measurements and analyses may be helpful in this.
- If so, should it be opt-in or opt-out?
 - For pragmatic reasons, we plea in favor of opt-out.
- What may CSIRTs expect from constituents in terms of receptiveness and feedback? And reversely, what may constituents expect from their CSIRT?
 - At a minimum, both parties should be able to reach the right person at the other party and via the right means of communication. Email may not be an adequate means for communicating urgent information. Therefore, a phone number and/or means of instant messaging is additionally recommended.
- How can a decision be reached within a short time frame (e.g., hours) on whether or not to perform proactive scanning for a particular new vulnerability?
 - Not every vulnerability can be tested as simply and reliably as CVE-2019-11510. This is the case when the attack vector to exploit a given vulnerability is undisclosed, when exploitation requires specific prior knowledge about the target, and when exploitation depends on a particular configuration. Depending on circumstances, timeliness may be more important than reliability of the initial test: gathering an initial set of *potentially* vulnerable systems, annotated as such, can provide a “worst-case” scope that can then be subject to further investigation.
- How should collected data (not) be processed, taking into account legal and ethical considerations including constituents’ autonomy, reputation, and privacy?
- Which IP ranges should (not) be included in the scans, and based on what rationale? Should constituent organizations provide the CSIRT with a list of IP addresses they use?
 - Traditional geographical boundaries may constitute legal/ethical boundaries. Unsolicited scanning that touches on services/equipment that is physically and/or logically associated with other countries might trigger a response from those countries’ CSIRTs or realms of cyber diplomacy. Note that a single constituent organization can have IP ranges and equipment in more than one country. Furthermore, the routing of scanning traffic can traverse geographical boundaries, even for domestic-to-domestic communication. That means domestic tests and the outcomes thereof may be observable/detectable by non-domestic parties.
- How to cope with IPv6, where subnets tend to be too big to be probed for asset discovery in full using existing methods and techniques?
- How to cope with intentionally deceptive infrastructure (e.g., honeypots) that triggers false positives, possibly tainting the outcomes of scans?

Because timeliness in identifying vulnerable systems can be critical, bureaucratic overhead should be kept at a minimum. A pragmatic attitude is desirable when confronted with vulnerabilities that can be exploited remotely without authentication. The questions listed previously acknowledge the work by Grobler and Bryk [4] in 2010

on the establishment of CSIRTs, for instance, regarding availability of resources and cultural and governmental differences. These aspects could also be studied in respect to longitudinal data collection for scientific purposes.

A last point of attention is information sharing: due to apparent differences in interpretations of the GDPR, the NIS Directive, and local laws and regulations, not all CSIRTs in the EU seem able or willing to share information about vulnerabilities in specific systems with parties outside their own constituency. As a result, information on identified vulnerable systems that was part of mass reports submitted to a national CSIRT did not always reach those who are in a position to act. This was observed in the Netherlands during the case of CVE-2019-11510¹⁰ and again in January 2020 during the case of CVE-2019-19781 (Citrix). One example of such differences is whether or not an IP address should be considered personal data in this context—even if the IP address is associated with an organization and not an individual person/consumer—and whether that should then prohibit dissemination of information. Legal and/or political decisions might be needed to enable governmental CSIRTs to coordinate dissemination of such information. That topic warrants attention but is not further discussed here. Interested readers are referred to the work of Cormack [2] and ENISA [3].

Meanwhile, the “elephant in the room” is that computer-indexing search engines operated by private companies, including but not limited to BinaryEdge, Censys, Shodan, and ZoomEye, already perform scanning at a global scale, collect large amounts of data, and make that data commercially available to anyone willing to pay for it. Today, Shodan displays possibly relevant CVEs next to data about systems and shows indices of unprotected ElasticSearch databases. Attackers make use of this open source information as well. That reality cannot be ignored.

ACKNOWLEDGMENTS

The author is grateful to the University of Amsterdam for providing support for this Field Note, and for input provided by Xander Bouwman, Jeroen van der Ham, Renate Verheijen, members of Dutch CSIRT communities (with a special mention for SURFcert), and the editors of DTRAP. The author welcomes feedback from readers.

REFERENCES

- [1] Kevin Beaumont. 2019. Pulse Secure SSL VPN Vulnerability Being Exploited in Wild. Retrieved January 20, 2020 from <https://opensecurity.global/forums/topic/184-pulse-secure-ssl-vpn-vulnerability-being-exploited-in-wild/>.
- [2] Andrew Cormack. 2016. Incident response: Protecting individual rights under the general data protection regulation. *SCRIPTed: A Journal of Law, Technology and Society* 13, 3 (Dec. 2016), 258–282. DOI : <https://doi.org/10.2966/scrip.130316.258>
- [3] ENISA. 2011. A flair for sharing—Encouraging Information Exchange Between CERTs. Retrieved January 26, 2020 from https://www.enisa.europa.eu/publications/legal-information-sharing-1/at_download/fullReport.
- [4] Marthie Grobler and Harri Bryk. 2010. Common challenges faced during the establishment of a CSIRT. In *Proceedings of the 2010 Information Security for South Africa Conference*. IEEE, Los Alamitos, CA, 1–6. DOI : <https://doi.org/10.1109/ISSA.2010.5588307>
- [5] NCSC-UK. 2019. Alert: Vulnerabilities Exploited in VPN Products Used Worldwide. Retrieved January 20, 2020 from <https://www.ncsc.gov.uk/news/alert-vpn-vulnerabilities>.
- [6] NSA. 2019. Cybersecurity Advisory: Mitigating Recent VPN Vulnerabilities. Retrieved January 20, 2020 from <https://media.defense.gov/2019/Oct/07/2002191601/-1/-1/0/CSA-MITIGATING-RECENT-VPN-VULNERABILITIES.PDF>.
- [7] Bad Packets. 2019. Over 14,500 Pulse Secure VPN Endpoints Vulnerable to CVE-2019-11510. Retrieved January 20, 2020 from <https://badpackets.net/over-14500-pulse-secure-vpn-endpoints-vulnerable-to-cve-2019-11510/>.
- [8] Pulse Secure. 2019. SA44101—2019-04: Out-of-Cycle Advisory: Multiple Vulnerabilities Resolved in Pulse Connect Secure/Pulse Policy Secure 9.0RX. Retrieved January 12, 2020 from https://kb.pulsesecure.net/articles/Pulse_Security_Advisories/SA44101.
- [9] tCell (since acquired by Rapid7). 2018. What’s Going on in Production Application Security 2018. Retrieved January 20, 2020 from <https://blog.rapid7.com/2018/08/22/whats-going-on-in-production-application-security-2018/>.

¹⁰One notable example is that some dozen organizations linked to the defense industry sector were among the hundreds in the mass report but were not contacted by the government. It was not until after the author himself contacted the Dutch Military Intelligence and Security Service (MIVD) that those parties were informed and had their systems patched or taken offline. In retrospect, this turned out to be due to the national CSIRT at that time being unaware of those particular IP addresses being in use by those particular organizations, as a result of which those IPs in the mass report were not recognized by the national CSIRT as belonging to defense industry organizations. This constitutes an important lesson learned.

- [10] Ars Technica. 2020. Unpatched VPN makes Travelex latest victim of “REvil” Ransomware. Retrieved January 20, 2020 from <https://arstechnica.com/information-technology/2020/01/unpatched-vpn-makes-travelex-latest-victim-of-revil-ransomware/>.
- [11] Orange Tsai and Meh Chang. 2019. Attacking SSL VPN—Part 3: The Golden Pulse Secure SSL VPN RCE Chain, with Twitter as Case Study! Retrieved January 12, 2020 from <https://devco.re/blog/2019/09/02/attacking-ssl-vpn-part-3-the-golden-Pulse-Secure-ssl-vpn-rce-chain-with-Twitter-as-case-study/>.
- [12] Jeroen van der Ham and Roland van Rijswijk-Deij. 2017. Ethics and Internet measurements. *Journal of Cyber Security and Mobility* 5, 4 (Oct. 2017), 287–308. DOI : <https://doi.org/10.13052/jcsm2245-1439.543>

Received October 2019; revised January 2020; accepted February 2020