On Federated Cyber Range Network Interconnection

Adamantini Peratikou[1], Constantinos Louca[1], Stavros Shiaeles[2] and Stavros Stavrou[1]

[1] Open University of Cyprus, Faculty of Pure and Applied Sciences, Nicosia 2220, Cyprus
[2] University of Portsmouth, School of Computing, Faculty of Technology, Buckingham Building, Lion Terrace, Portsmouth PO1 3HF, United Kingdom

Abstract. Cyber Ranges exist to enable hands on training within realistic ICT infrastructures in a sandboxed environment, to investigate attack and defense strategies and to assess the resilience of the infrastructures. To fully exploit their capabilities one has to have access to multi domain exercises, which may combine ICT, naval, electrical grid, telecom or other relevant infrastructures. It can become obvious that no single organization can easily own or sustain a multi domain cyber range and that there is a need to connect multi domain Cyber Ranges from different organizations together. This paper focuses into analyzing the current state of the art on the federation of Cyber Ranges, by focusing on the federated network interconnection. Various methods for interconnecting distributed Cyber Ranges into a single federated Cyber Range are being discussed and their network performance impact is evaluated. VPNs are widely used to interconnect networks together due to their relative low cost and simplistic nature, however, performance of the network must be accounted, alongside the flexibility the VPNs can provide to support multiple scenarios in a multi domain distributed federated Cyber Range. This work focuses on the performance comparison of IPsec and Virtual Tunnels.

Keywords: Cyber Ranges, Federation, Interconnection, VPN, IPsec, ,OpenVPN, Virtual Tunnels, Cyber Security

1 Introduction

In an increasingly networked connected world, where successful cyber attacks can have disastrous effects, which are not retained to the targeted system but can also affect multiple sectors of the economy, it is of crucial importance that cyber security practitioners are trained in realistic multi domain training environments. This would allow them to hone their skills and enable them to have cyber attack response experience without actually being exposed to a real attack.[1] Given that it is inadvisable to use a live system for training purposes, or to take it offline so that users can train on it, the need arises for creating Cyber Ranges (CRs). CRs of this kind tend to be focused on that area of expertise specific to the entities running them, and while they can be quite comprehensive in the study of their own subsection of cyberspace, they are often limited to that subsection [2].

For example, while a CR focusing on testing the vulnerability of an e banking system can be very helpful in allowing researchers to draw conclusions on securing e banking systems, it can offer no insights as to the possible consequences of a particular threat to another system, for example Electrical Grid oversight, or Air Traffic Control [3]. This is to a large extent due to the fact that the threats to one type of system are erroneously perceived as unrelated or irrelevant to another, as well as to the fact that testing for the effects of these threats on a system requires very specific information on the functioning of that particular system. Different systems have different services, resources and requirements, and in order to draw accurate conclusions, one needs to be able to study a threat across a wider range of systems, including some that one would not normally have access to, or expertise on.
The Federation of CRs attempts to fill this gap, by pooling together different CRs form similar or different domains, and testing for the effects of a particular threat across a more widespread selection of targeted systems, emulating therefore the true nature of the interconnected world. This paper examines both the theoretical soundness of pooling resources in this manner, as well the technical methodology required to do so.

The first section of this paper provides some possible techniques of implementing the interconnection of federated CRs by overviewing different technologies. The second section presents relevant performance metrics. The third section analyzes the requirements and proposes an interconnection architecture to satisfy those requirements. The final section provides some preliminary results obtained via performance tests to identify the appropriateness of the solution.

2 Related Works

The European Defense Agency currently works on a project for creating a federated CR environment [2] where EU Member States will federate their national CRs and improve their respective Cyber Defense training capabilities. The congressional record proceedings of the 113th congress session of the United States [4], documents an attempt to interconnect individual CRs together in a distributed manner. Another example is the EU funded project ECHO and the CyberSec4Europe that upon completion will analyze requirements, specifications, and use cases for federation of CRs [5]. Also, the EU funded project Foresight aims to provide a multi domain geographically distributed Federated Cyber Range [6]. A number of CRs are currently in use or being developed and provided by third party organizations across governments, academia and private sector. CRs provided by universities are predominantly used for education and training, most of them have just emerged in the last few years. Currently the most established CR in a university setting is in Michigan and is being governed by 12 public universities, named the Merit Network [7]. The CR has been available since 2012, it has four physical locations and incorporates a virtual training environment called “Alphaville” to evaluate cybersecurity skills. Another example is the CR called “The DETERLab” at the University of Southern California [7]. CRs provided by industry are primarily used for the training of cyber security

professionals or used from other companies for finding vulnerabilities in their own infrastructures [8].

While several governments are investing in CRs, and research is currently conducted to interconnect CRs together, the information that is publicly available is limited. Substantial research interest has been dedicated to technologies that enable the interconnection of CRs, but there is a significant gap in the literature when it comes on interconnecting CRs with the use of VPNs.

Various VPN technologies exist that are based on Point to Point Tunneling Protocol (PPTP), IP Security standard protocol (IPsec) or SSL (Secure Sockets Layer) technology [9]. PPTP are not generally preferred due to security issues, arising from the simplicity of the protocol [10]. For interconnecting CRs in a federated environment, IPsec and SSL/TLS based VPN tunnels will be compared. IPsec and SSL are a set of cryptographic protocols that provide secure communication. IPsec protocols include AH, ESP, IKE, ISAKMP/Oakley, and transforms [11]. IPsec connection starts with a two phase handshake and when it is completed, arbitrary traffic can be sent via an encrypted tunnel [12].

By decoding the current trend and the need for realistic training and evaluation, we can foresee more and more CRs being developed in the forthcoming years, where the individual CR capabilities can be extended by interconnecting individual CRs together in a federated environment.

3 CR Interconnection Framework

3.1 Framework metrics

A federated CR interconnection environment above all must be reliable and efficient. To assess the effectiveness of such environment, or framework, a set of measurements need to be satisfied. Various metrics span across research communities and agencies tasked with cybersecurity [13]. A general framework without any detailed recommendations, ITIL, is used by organizations for service quality management, while COBIT is used by several organizations as a measurement technique in security areas. Considering that the federation of CRs is still in an evolving state, we are proposing a set of minimum measures that need to be satisfied.

Scalability: Due to the emerging nature of CRs the interconnection framework should be scalable enough to support an increasing number of CRs to be connected in the federated environment.
Fault tolerance: The interconnection of the federated CRs has to be resilient and able to intelligently handle faults as the VMs under the CRs toggle between offline and online states and participate in exercises across geographically distributed multi domain CRs.
Availability: Availability is defined as the probability of receiving a fulfilled service request in an acceptable time [14] The interconnection of CRs has to be up at all times and provide the service requested by CR users within a reasonable time.

Performance: For the performance metrics, a set of attributes must be taken into consideration such as Response time, throughput, CPU load, memory and network

usage. Also depending on the interconnection technology, whether VPN or IPsec, the number of simultaneous users that can be supported is also required.

Security: Security can be interpreted in terms of confidentiality, integrity and availability which are considered difficult to measure but are assumed that are handled by a well accepted secure protocol like IPsec etc.

3.2 CR Federation Requirements

A realistic federated Cyber Range training environment should allow the interconnection of any CR domain environment through its network architecture. This capability enables an increased number of end users, belonging to different domains, to be trained in a multi domain environment. By implementing federated gateways, and relevant functionality, a realistic federated Cyber Training environment between remotely distributed platforms can be established. Federated gateways can be realised through the use of Virtual Private Networks (VPNs). The addressing schemes in such methods must be agreed upon the CRs so that each CR has the appropriate address space in the subnets that will be used in the distributed federated scenarios.

It is envisioned that organisations that own a CR or cyber training environment, will use what we call a federated gateway, to connect their own platform or resources to other CRs running their own federated gateway. This interconnection will allow the combination and sharing of remote resources, to form a very large scale distributed federated CR.

3.3 Interconnection Requirements

The interconnection architecture considers the various cyber ranges and connects them using a centralized federated gateway to form a unified platform. The architecture should be designed in such a way to allow the provision of management, deployment, and usage of VMs and Cyber scenarios of interest. With that in mind the remote gateways should have a centralized master federation point of reference to serve as the management for the whole federation platform environment e.g. to monitor and manage the status of the interconnection links at any given time and take corrective action in case of a CR link failure.

The requirements of the federated gateways are as follows:

- To support ethernet TCP/IP communication
- Federated gateways will be based on VPN technologies
- Point to multipoint, or mesh connection
- Monitoring service/process to constantly monitor connectivity between CRs

3.4 Interconnection Techniques

A federated interconnection can be implemented at layer 1, layer 2 or layer 3. Interconnecting at layer 1 is very challenging for supporting complex scenarios and exercises and has practical cost and implementation implications. On the other hand,

layer 2 and layer 3 interconnectivity are less costly to establish and have the potential to grow to an ad hoc federation of geographically distributed interconnected CRs across the world. Virtual Private Networks (VPNs) allow secure connections to a private network, across any public network, and replace the need to physically lay private cables over long distances [15]. The most appropriate method to allow federation between CRs is with the use of VPNs due to their proven cost effectiveness.

For the accomplishment of layer 2 VPN services, Figure 1, the traffic on layer 2 VPN is transported to the service provider network by MPLS, which on the sending and receiving ends it gets transformed back to Layer 2 format. Different Layer 2 formats can be configured at the sending and receiving ends. The security and privacy of an MPLS Layer 2 VPN are equal to those of an ATM or Frame Relay VPN [16]. The service provisioned with Layer 2 VPNs is also known as Virtual Private Wire Service (VPWS). Generally, on a layer 2 VPN the routing is done on the customers routers, therefore CR providers routers. Those routers must choose the correct circuit to send the traffic to. In the provider edge, the routers receiving the traffic send it across the service provider network. For a Layer 2 VPN, the clients routers need to be configured to carry all layer 3 traffic, while the routers on the service provider only require to know the amount of traffic the VPN layer 2 will carry, to carry that traffic between the client using layer 2 VPN interfaces [17]. The policies must be configured on the provider routers. The clients are only required to know which VPN interfaces connect to which of their own sites. Figure 1 illustrates a Layer 2 VPN in which each site has a VPN interface linked to each of the other client sites.

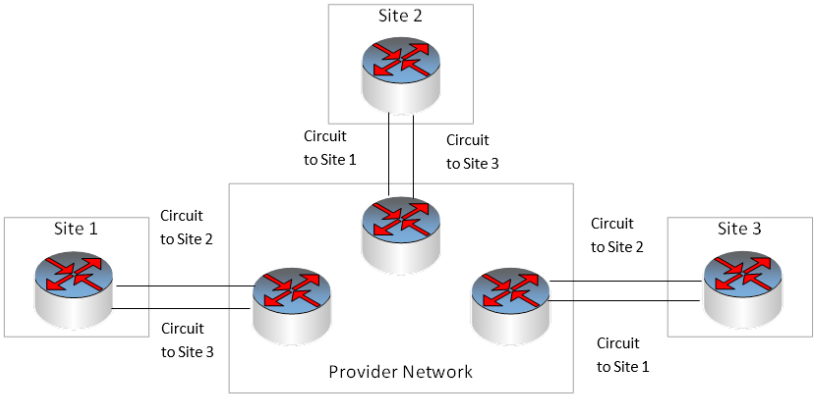

Fig. 1. Layer 2 VPN Connecting Routers

While layer 2 VPN can connect CRs in one big federation environment, there would be a need for additional control mechanisms for monitoring and managing those CRs. While planning this kind of connection, an input from every range architecture is needed. Different IP addressing and VLAN schemes, environment isolations and implementation of monitoring and situational awareness tools must be agreed among the interconnected CRs. To avoid this complexity and provide a more controlled CR

environment, where each CR has more security control over their own interconnected CR, it was decided that a layer 3 VPN connection will be implemented and tested for the federated gateways. The CRs need to exchange resources with the use of site to site VPNs. The options for layer 3 site to site VPNs are IPsec and SSL/TLS from OpenVPN. IPsec establishes VPN tunnels through the use of protocol and cryptographic security measures. It supports network level peer authentication, access control, connectionless integrity, data origin authentication, detection and rejection of replays, confidentiality, and limited traffic flow confidentiality [18]. IPsec uses two modes of operation, the tunnel mode and the transport mode. Tunnel mode is used in site to site communication and it encrypts the header and the payload. The protocols utilized by IPsec to allow traffic security are the Encapsulating Security Payload (ESP) and the Authentication Header (AH). ESP offers integrity data origin authentication and confidentiality. AH is used for packet integrity validation and sender authentication. IPsec connects peers through the Internet Key Exchange (IKE) protocol [19]

OpenVPN utilizes the SSL/TLS protocol for cryptographic elements required by VPN and tunnel creation [20]. OpenVPN instead of accessing the network interface like IPsec, it generates instead a virtual network interface. It encrypts traffic using the same principles for the handshake as IPsec IKE, and SSL libraries are then used to secure the symmetric tunnel.

3.5 Network Architecture

The high level network architecture of a federated CR is depicted in Figure 2.
This federated network serves as a virtual environment where all multi domain cyber security scenarios can be executed. The number of CRs can increase dynamically, in case a new CR vendor is interested in joining the federation.

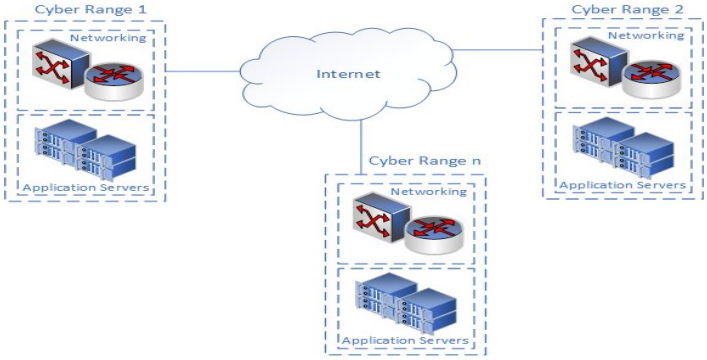

Fig. 2. Architecture of the federated CR network

The proposed interconnection architecture utilizes Layer 3 VPN technologies. VPN technologies that can be used in a federated CR environment consist of the standard remote VPN that will be responsible for routing all the traffic using specific routing policies and site to site virtual tunnels. A site to site setup will allow different networks to be connected using the VPN tunnel. To promote collaboration and accessibility, open

source packages are used on federated gateways. An OpenVPN client/server will be used to setup an encrypted remote VPN.

In the point to multipoint configuration, the federated gateway controllers require a host to act as the central federated controller (Figure 3). Each additional CR location will be set up with a connection profile connected to the central federated controller. Once the server is configured accordingly and the connection profile is installed on the CR's hosting machine, the CR will be connected to the federated network. Then, CRs will be automatically connected to all the resources that the central federated controller provides. This means that the server security controls can be enforced site to site, and the resources of the network are accessible at all locations. IPsec site to site requires a specific set of ports, 400 and 4500, to be available by all participating CRs while OpenVPN site to site tunnels requires one port per tunnel [21].

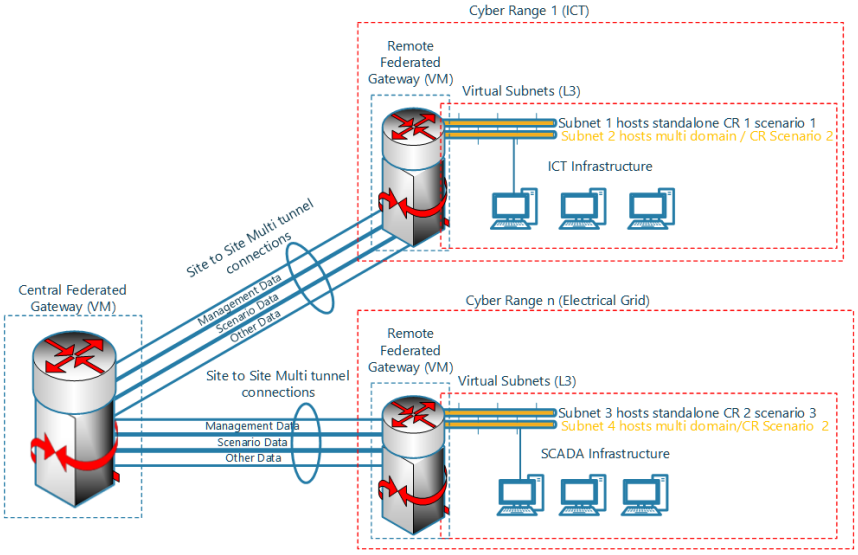

Fig. 3. Proposed Interconnection

Monitoring

A process is required to monitor the federated gateways' controller status of the various CRs at all times. To achieve this a web based application was designed. The application displays the status of each CR that serves as a gateway controller (server), including all current connections (clients). The application works by parsing VPN data from the controller logs. The web based application shows all the relevant information for each VPN server and client. For the server it highlights the VPN mode, status whether connected or disconnected, number of clients, traffic, UP time and local IP of the DHCP server on the federated gateway. For the connected client it illustrates the hostname, the VPN IP & Remote IP, the location of the client, the traffic, the last contacted (last ping) and the uptime. A map is also generated to show the exact location of the deployed CR. The web application can also serve as a simple heartbeat monitor for all gateway clients and server instances.

4 Results

To characterize the network performance of the different CR interconnecting technologies, iperf3 was used. A basic test setup is illustrated in Figure 4. The iperf3 tool was run through a script that was automatically executed on a predefined time interval. The script collects the IPs of the connected federated gateway controllers (CRs) from the monitoring tool, specified above, and runs iperf3 TCP and UDP tests on each. Because the most common use of iperf3 is to measure the speed of data transfer, the reported results will be the bytes per second of the data transfer rather than the Layer 3 network performance transfer characteristics. To compensate for this, a set of multipliers are being used to include the calculated overhead for specific packet lengths and to convert from iperf3 results to Layer 3 and Layer 2 measurements[22]. Iperf3 is configured by default to take advantage of large packets and TCP window scaling. For testing, specific options were used to force different package sizes. To achieve the packages sizes iperf3 requires configuring flags such as the length option (l) and the TCP MSS (M).The length flag represents the payload (minus TCP timestamps) and the MSS option represents the payload and 12 bytes for the TCP timestamp option.

Fig. 4. Performance tool architecture

4.1 Tests parameters

Throughput tests were run on both VPN technologies and over an unencrypted connection under various packet sizes.
Tests were run over a 1Gbit/sec ethernet connection. Testing environment includes, two virtualized Debian based routers, one acting as a central federated controller that resides on one CR, and a second instance residing on a different CR, and two virtual machines (VMs). A Debian VM on the first CR site connected to a subnet of the first router (RT1) and an Ubuntu VM the second site connected to a subnet of the second router (RT2). The Debian based routers are running VyOS (a Linux based network router and firewall distribution) [23]. The network layout and environment infrastructure of IPsec and OpenVPN virtual tunnels were identical, as presented in Figure 3.

Figure 5 illustrates the TCP throughput obtained when a packet size of 128 1024 is set for both VPN configurations. The VM residing on RT1 subnet acted as a server and the second VM on RT2 subnet as a client as presented in Figure 4. The throughput is normalized, expressed as the ratio of the throughput obtained from the VPN technology

under test, to the throughput of the same unencrypted connection. As expected, the throughput increases proportionally with the packet load. The results show that the throughput of IPsec surpasses by up to 8% the throughput of OpenVPN virtual tunnel for all packet sizes. Figure 6 shows the latency for IPsec, OpenVPN and the unencrypted ethernet connection. IPsec has slightly lower latency than OpenVPN virtual tunnels. Latency for both VPN protocols, as expected, is noticeably higher than the unencrypted ethernet configuration. Results presented in Figure 5 and Figure 6 suggests that the performance of the two protocols is comparable.

For the same scenario as above, UDP throughput tests were also conducted, shown in Figure 7. The performance of the two VPN technologies was found to be comparable for packet sizes 128 256, while for packet sizes 512 1024 IPsec performance marginally decreases. The difference between the two VPN configurations is too minimal to be considered as an advantage for one versus the other.

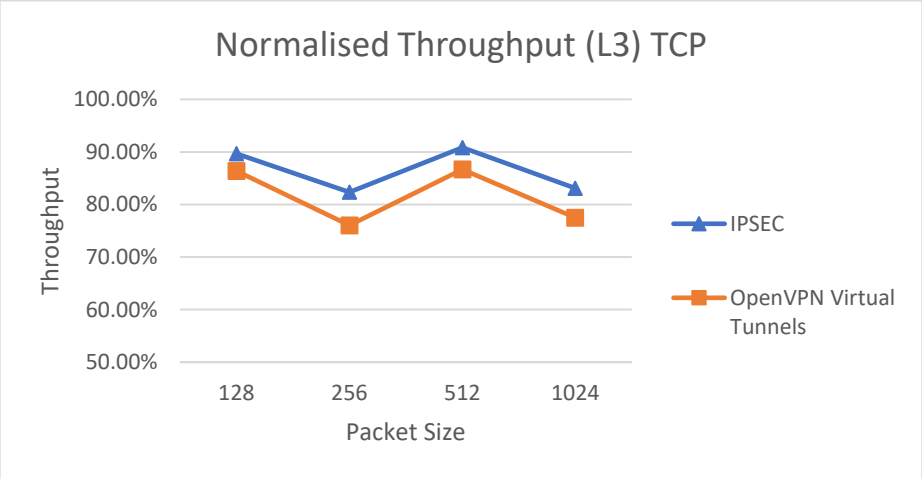

Fig. 5. Normalized Layer 3 TCP Throughput (a VM residing on Router 2 (Client) to a VM residing on Router 1 (Server))

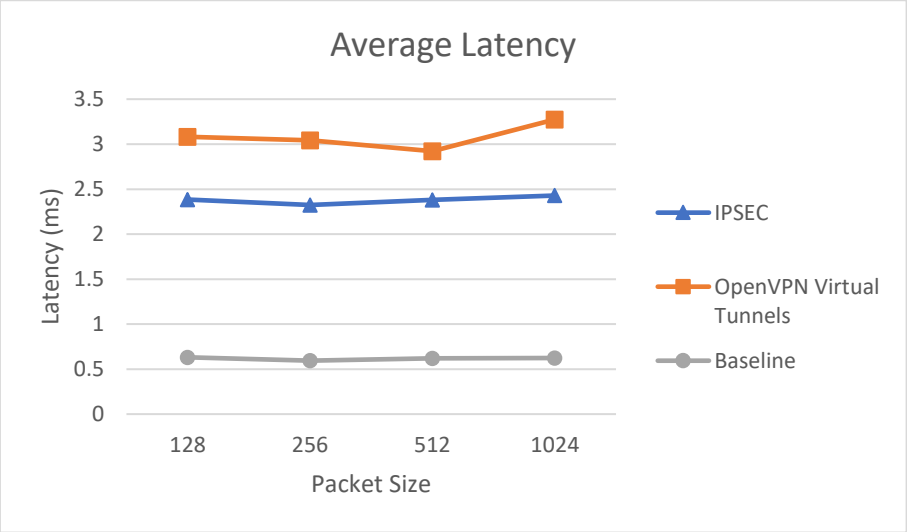

Fig. 6. Latency (VM residing on Router 2 (Client) to a VM residing on Router 1 (Server))

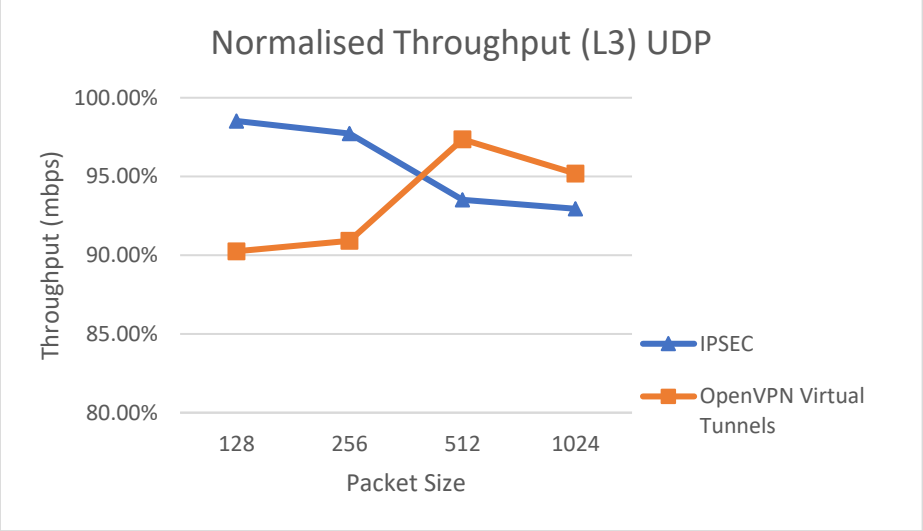

Fig. 7. Layer 3 UDP Throughput (a VM residing on Router 2 (Client) to a VM residing on Router 1 (Server))

5 Conclusions

This paper overviewed and compared possible techniques for realizing a generalised interconnection architecture for federated CRs. The framework performance metrics that need to be taken into account when interconnecting CRs have

been presented. Preliminary results obtained through network tests to identify the performance of the protocols involved have been also provided.

Our evaluation encompassed a wide range of packet sizes and identified that IPsec and OpenVPN virtual tunnels behave similarly in all packet sizes with no major differences when it comes to performance. In general, the performance depends on the interconnecting link bandwidth, the speed of the encrypting and decrypting device and the type of cipher in use. The proposed federated interconnection architecture illustrated in Figure 3 can be implemented using either OpenVPN or IPsec depending on the availability of network ports, with respect to how many CRs are involved in the federation and how many geographically distributed subnets of these CRs are involved in the scenarios under investigation.

Acknowledgement

The authors would like to acknowledge the FORESIGHT project funded by the European Union's Horizon 2020 research and innovation programme (grant agreement: 833673), and the partners on the project.

References

1. Debatty, T., Mees, W.: Building a CR for training CyberDefense Situation Awareness. In: 2019 International Conference on Military Communications and Information Systems, ICMCIS, pp. 1 6, Budva, Montenegro, doi: 10.1109/ICMCIS.2019.8842802 (2019).
2. Ferguson B., Tall A., and Olsen D.: National Cyber Range Overview. In: 2014 IEEE Military Communications Conference. IEEE, pp. 123 128. doi: 10.1109/ MILCOM.2014.27 (2014).
3. Ellis, R., Mohan, V.: Rewired Cybersecurity Governance . John Wiley & Sons Incorporated, Somerset (2019).
4. Congressional Records: Proceedings and debates of the 112th congress. United stated government publishing office, Washington DC (2012).
5. Directorate General for Internal Policies: Cybersecurity in the European Union and Beyond: Exploring the Threats and Policy Responses, European Parliament (2015).
6. FORESIGHT Project Homepage, https://foresight h2020.eu, last accessed 2020/07/10.
7. Urias, V., Stout, W. M.S., Van Leeuwen, B., Lin, H.: CR Infrastructure Limitations and Needs of Tomorrow: A Position Paper, United States. doi: 10.1109/CCST.2018.8585460 (2018).
8. IBM, IBM Invests $200M, https://www.enterprisetimes.co.uk/2016/11/17/ibm spends 200m cyber range/, last accessed 2020/07/12
9. Berger, T.: Analysis of current VPN technologies. In: 1st International Conference on Availability, Reliability and Security (ARES'06), IEEE, pp. 108 115, Vienna, doi: 10.1109/ARES.2006.30. (2006).
10. Kotuliak, P.R, Trúchly P.: Performance comparison of IPsec and TLS based VPN technologies. In: 9th International Conference on Emerging eLearning Technologies and Applications (ICETA), pp. 217 221, Stara Lesna, doi: 10.1109/ICETA.2011.6112567 (2011).

11. Dhall, H., Dhall, D., Batra, S., Rani, P.: Implementation of IPsec Protocol In: 2nd International Conference on Advanced Computing & Communication Technologies 2011 , pp. 176 182, Rohtak, Haryana, doi: 10.1109/ACCT.2012.64.I (2012).
12. Schneier, B., Mudge : Cryptanalysis of Microsoft's Point to Point Tunneling Protocol (PPTP). In: Proceedings of the 5th ACM Conference on Communications and Computer Security, pp. 132 141, ACM Press (1998).
13. ENISA: Measurement Frameworks and Metrics for Resilient Networks and Services: Challenges and Recommendation. European Network and Information Security Agency (ENISA) (2010).
14. Monakhov, Y. M., Monakhov, M.Y, Luchinkin, S. D., Kuznetsova, A.P., Monakhova, M.M.: Availability as a Metric for Region Scale Telecommunication Designs. In: 10th IEEE International Conference on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS), pp. 775 779, Metz, France, doi: 10.1109/IDAACS.2019.8924390 (2019).
15. Narayan, S., Williams, C. J., Hart, D. K., Qualtrough M. W.,: Network performance comparison of VPN protocols on wired and wireless networks. In: 2015 International Conference on Computer Communication and Informatics (ICCCI), pp. 1 7, Coimbatore, 2015, doi: 10.1109/ICCCI.2015.7218077 (2015).
16. Kompella et al.: Layer 2 Virtual Private Networks Using BGP for Auto Discovery and Signaling. In: IETF RFC, May 2012, https://tools.ietf.org/html/rfc6624 (2012)
17. Metz, C.: The latest in VPNs, part II. In: IEEE Internet Computing, vol. 8, no. 3, pp. 60 65, May June 2004, doi: 10.1109/MIC.2004.1297275 (2004).
18. Kent, S., Seo, K.: Security Architecture for the Internet Protocol (No. 4301). In: IETF RFC, December 2005: http://www.ietf.org/rfc/rfc430l.txt, (2005)
19. Hauser, F., Häberle, M., Schmidt, M., Menth, M.,: P4 IPsec: Site to Site and Host to Site VPN with IPsec in P4 Based SDN. In: IEEE Access, doi: 10.1109/ACCESS.2020.3012738 (2020).
20. Qu, J., Li, T., Dang, F.,: Performance Evaluation and Analysis of OpenVPN on Android, In: Fourth International Conference on Computational and Information Sciences, pp. 1088 109, Chongqing, doi: 10.1109/ICCIS.2012.203 (2012).
21. Kivinen, et al.,: Negotiation of NAT Traversal in the IKE(No. 3947). In: IETF RFC, January 2005 https://tools.ietf.org/html/rfc3947 (2005)
22. Soucy, R.,: Network Router Performance Testing How To., http://soucy.org/vyos/NetworkPerformanceTesting.pdf, last accessed 2020/08/02.
23. Vyos Homepage, https://www.vyos.io/, last accessed 2020/07/13.