Managing Risks in Enterprise VPNs: A Framework

MSc Research Project

MSc Top-Up Cybersecurity

Darragh Gavin

Student ID: 22157468

School of Computing

National College of Ireland

Supervisor: Ross Spelman

Link for VIVA

National College of Ireland

MSc Project Submission Sheet

School of Computing

Student Name: Darragh Gavin………………………………………………………………………………

Student ID: 22157468………………………………………………………………………………………..…

Programme: MSc Top-up CyberSecurity…………………… **Year:** 2024………………..

Module: MSc Research Project……………………………………………………………….……

Supervisor: Ross Spellman……………………………………………………………………….………

Submission Due Date: August 12th …………………………………………………………………….……

Project Title: Managing Risk in Enterprise VPNs: A Framework………………………….……

Word Count: 5964…………………………… **Page Count**…17……………………………………..

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.
ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature: ………………………………………………………………………………………………………

Date: ………………………………………………………………………………………………………

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	□
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	□
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	□

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Managing Risk in Enterprise VPNs:
A Framework

Darragh Gavin
22157468

Abstract – With the widespread adoption of Virtual Private Networks during COVID, VPNs proved a valuable target for hackers looking for entry into enterprise networks. Attacks have risen dramatically in recent years, and further investigation into the vulnerabilities exploited seemed warranted. This is done with a view to categorising the attacks and building a model that outlines appropriate steps at remediation and limiting of risk where possible. If codified into a framework, this should provide a path to greater overall security while allowing for remote access to networks.

Index Terms – VPN, Cross-protocol attack, tunnelling, IPSec, OSPF, DoS

1 Introduction

The internet is made possible by connecting private servers and networks via public (or publicly accessible) infrastructure, allowing data to travel between endpoints. When we log onto a platform or visit a site, it feels as though we are connecting directly. Realistically, this resource is probably held on a server in a completely different country and our communication traverses a web of infrastructure to enable this connection. In most cases, we use secure connections, such as HTTPS, when connecting to public domains. This approach isn't sufficient for corporate environments, where staff are accessing their corporate network.

This is where Virtual Private Networks provide an invaluable service, in enabling devices from outside the network to behave as though they were directly connected and securely communicate. This is done using groups of protocols that allow for secure encryption, authentication, and data integrity.

The global pandemic brought security into focus, and with 93%[1][2] of organisations now using VPNs vulnerabilities in VPN infrastructure warrant further investigation. Remote working is now a staple in many industries and, while the use of VPNs is down from its peak, this is still a threat vector that is causing problems for many companies. In Threatlabz[3] recently published VPN security report, 56% of cybersecurity professionals reported having cyberattacks related to their VPN in the previous 12 months. The report showed that 35% experienced more than one attack. This data was gathered from 600+ security professionals and dealt with the expanding threat surface provided by VPN usage.

[1] https://dataprot.net/statistics/vpn-statistics/

[2] https://www.statista.com/statistics/1343692/worldwide-virtual-private-network-reasons-usage/

[3] https://zerotrust.cio.com/wp-content/uploads/sites/64/2024/05/threatlabz-vpn-risk-report-2024.pdf

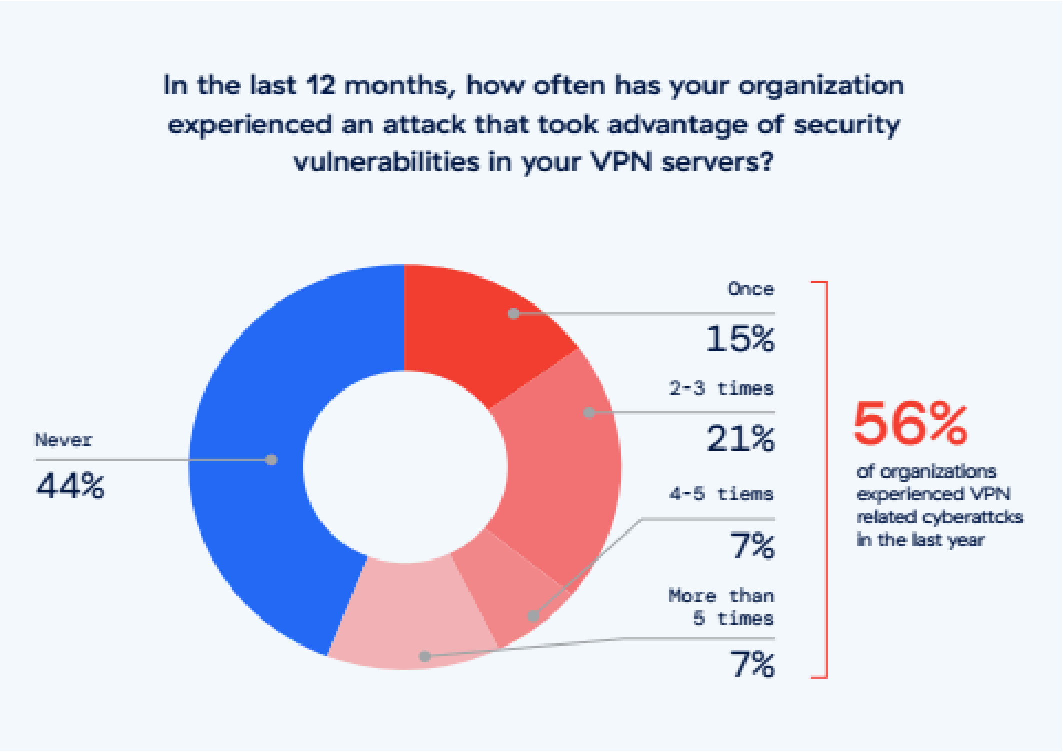

Diagram from Threatlabz VPN Risk Report 2024 – link in footnotes

The goal of this paper is to examine the risks facing organisations deploying VPNs to enable remote communication, whether to facilitate remote working or secure connections between remote facilities. When we look at the threat landscape for VPNs and the areas targeted over the last decade, patterns emerge in which areas successfully exploited. VPNs require Confidentiality and Integrity, which requires encryption, communication protocols, infrastructure, and then client/server applications to manage these connections. It stands to reason that there are avenues worth pursuing that are harder to secure.

In preparing this report, a wide range of materials were used including research, documentation, attack reports, CVE records, and published literature. This was then used to build a framework with one aim - ascertain best practices in mitigating vulnerabilities, and then codify these best practices reducing the risk associated with enterprise VPN usage. This was modelled in a lab environment to show its application and (hopefully) its efficacy.

1.1 Background

VPNs are an attempt to overlay private networks on top of publicly accessible infrastructure. This creates what feels like a direct link between gateways or endpoints, that provides secure communications – a virtual private network. There are generally 3 steps involved in creating this secure channel (key exchange, agreement, and session) and these steps have the potential for compromise, though we will see later that some are more vulnerable than others.

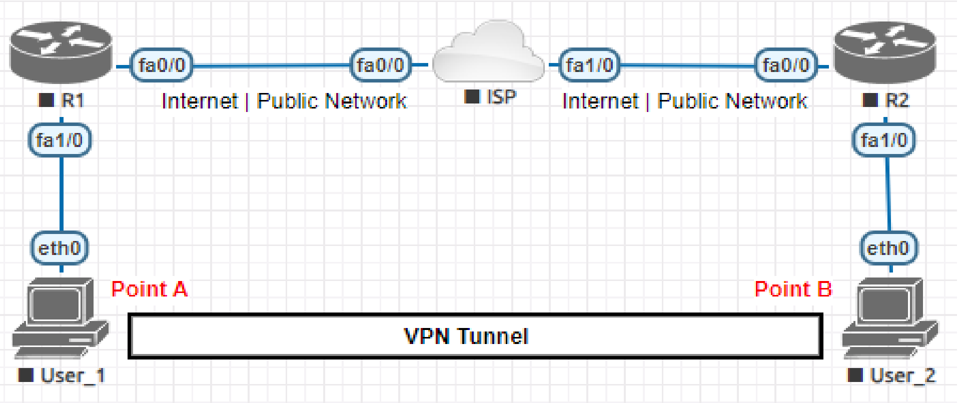

VPN connection, indicative of Site-to-Site VPNs

There are 3 primary enterprise VPNs, though MPLS is generally used by service providers and outside the scope of this paper. This leaves us with site-to-site (shown above) and Remote Client - usually software based and can be used by 3rd party vendors (shown below).

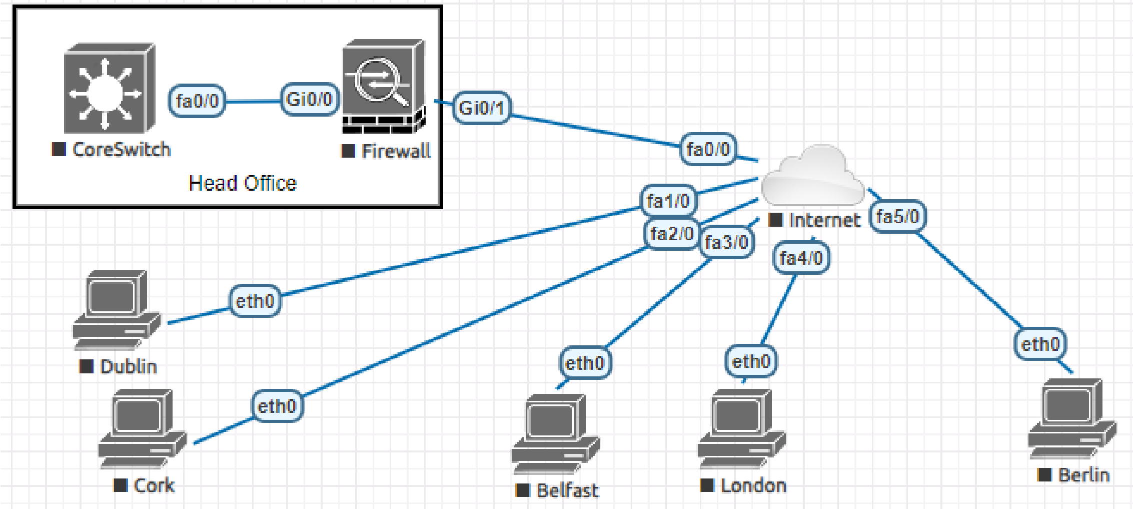

Client-based VPN Connection Model

1.2 VPN Protocols

There are many VPN protocols currently active and available, with IPSec and OpenVPN being the most commonly used in commercial settings. IPsec (Internet Protocol Security) is a group of freely available (open) protocols that was developed by the IETF. Not long after, a group led by a Microsoft produced another VPN protocol, called PPTP (Point to Point Tunnelling Protocol). Cisco later developed L2F (Layer Two Forwarding) which the IETF combined with PPTP to produce L2TP. While PPTP is no longer in use, L2TP is still combined with other protocols (such as IPSec) for VPN tunnelling.

There are 4 key aspects to VPN protocols (outside of integrity and confidentiality) which are

- Encapsulation
- Key exchange
- Encryption
- Tunnelling

When we talk of encapsulating in the context of network traffic, we are talking about specific layers in the TCP/IP stack (shown below) being housed within a package at a lower layer. Each layer is responsible for the inclusion of specific data. In the case of VPN tunnelling, the encapsulation happens at a higher layer or, to put in more succinctly, VPN tunnelling is encapsulating data from one protocol within the data of a protocol at the same or higher layer (Snader, 2006).

The key exchange protocol is very relevant, as it has been prone to issues when older protocols are used. Encryption standards are also very important though have changed less in recent years. There was a huge jump in computing power in the late 90's that necessitated a complete revision in the strength and complexity of encryption (which was itself enabled by the updated computing power). With the relatively recent inclusion of elliptic curves and the advent of quantum computing, it's still very much a race to ensure encryption remains sufficiently robust.

The process of tunnelling depends on the gateways (or endpoints) used to create the tunnel. In site-to-site VPNs, when we talk of tunnelling, usually we think of layer 3 (network layer) or layer 4 (transport layer) protocols, such as GRE, SSTP, or IP in IP. There are also layer 2 protocols (data link layer) like L2TP. Understanding this is central to understanding VPNs, as different security protocols are applied to communications at different layers and are bundled to create different VPN protocols (Jahan, et al., 2017).

OSI Layer	TCP/IP	Tunnelling/VPN Protocols
L7 – Away	Application Layer	SSL TLS DTLS
L6 – Presentation		
L5 – Session		
L4 – Transport	Transport Layer	OpenVPN | VXLAN | SSTP | SSH
L3 – Network	Internet Layer	IP in IP | GRE | IPsec
L2 – Data (Link)	Link Layer	L2TP PPTP
L1 – Physical		

Table of Tunnelling and VPN Protocols

2 Related Work

In researching this topic, I divided the available literature into two categories – namely VPN implementation and VPN security. The former looks at the progression of VPN architectures and best practices when setting up site-site VPNs or a client-server model. The latter looks at specific security postures and how they relate to attacks on VPNs, as well as the attacks themselves. To some extent, these could be seen as two sides of the same coin, but in practice, there are discrepancies that point to a need for a clear framework that helps align best practices with commonly overlooked vulnerabilities.

That is the object of this paper – to outline a clear framework that minimises/optimises organisational risks stemming from increased reliance on VPNs.

2.1 VPN Implementation

In discussing best practices for VPN deployment, there tended to be two perspectives taken. The first was a unified approach that seemed closer to a "defence in depth" model, as seen by (Pedapudi & Vadlamani, 2022). They discuss key areas to be secured to provide a functional VPN, including routers, gateways and firewalls, access policies and profiles, as well as working towards a unified threat management system, such as next gen firewalls that provide IPD/IDS. This is taken a step further in Tongkaw's case study which includes 13 "categories" that each have their own security overhead and requirements (Tongkaw, 2021). The categories cover the different security postures of 6 universities in Thailand, and how they mitigate, or fail to mitigate, a variety of architecture and policy issues. This included hardware used, mode of access, end device security, and monitoring.

In my opinion, this is a fantastic starting point but not a comprehensive implementation as it overlooks other aspects of VPN security, such as client software or 3rd party access.

This is further highlighted in a recent paper, where network configuration and management are the primary culprits in vulnerable VPN services (Tay, et al., 2022). The three primary problems listed are unmanaged access, old protocols, and firewall misconfiguration. Again, the recurring theme of oversight is present, but is insufficient in addressing a wider range of potential issues.

The second approach is one that looks at the core technologies used in creating a VPN (Xu & Ni, 2020), and primarily covers: Tunnelling, Encryption, Authentication, and Key Exchanges. This second approach also includes VPN specific issues, such as split tunnelling (Rao, et al., 2023) and latency (Jahan, et al., 2017), and their impact on network performance and security.

It's also worth directly mentioning the security assessment survey which provides quite an extensive review of the research into VPN performance, protocols, and attacks (Abbas, et al., 2023). Their paper falls short in its engagement with the materials and, to some extent like this paper, it covers too broad a subject matter. It doesn't look to draw any conclusions or propose solutions, but it does demonstrate the difficulties with topic that encompasses so many technologies. Of note was its regular reference to fundamentals like password policies and protocol selection, as well as Uskov's research on VPN selection strategies which generally favoured IPsec for enterprise VPNs (Uskov, 2012).

2.2 VPN Security and IKEv1

One issue that came up repeatedly was the vulnerabilities in IKEv1, which was compounded by a number of issues with operating systems (e.g. LibreSwan, CVE-2023-38712) and allowed for crashes and Denial-of-Service conditions. In (Sawalmeh, et al., 2021), we see a theme of network access via insufficiently secured nodes that allows a Denial of Service on the VPN gateway by flooding it with UDP packets (e.g. Cisco, CVE-2024-20308). At first glance, it looked highly unlikely that this kind of access would be given, but many public facing organisations are highly accessible. The second part of the attack centres around IKEv1's handling of ISAKMP packets. The recommendation was to protect the edge node with a passive configuration so that it ignores LSA's, but I think that falls short. A more comprehensive response that includes an upgrade to IKEv2 is required.

One key issue that is ubiquitous in security discussions is the overhead incurred with stronger encryption standards. In analysing the trade-offs made between security and throughput (Mahmmod, et al., 2020), they argue that a new approach to ensuring data authentication and confidentiality is needed. In order to reduce the overhead associated with secure encryption standards, they suggest using Field Programmable Gate Array chips, which would allow for parallel processing of encryption algorithms while handling traffic. In framing this solution, the authors note that it isn't just about improving network performance, but also enabling stronger security.

While I agree with this sentiment, I think we have sufficient computing resources to apply adequate encryption, particularly given the adoption of elliptic curve cryptography in the Diffie-Helmann key exchange.

2.3 TLS

Another issue, relating to protocols, was with SSL/TLS. There are two attacks that are synonymous with cross-protocol attacks, namely ALPACA and DROWN (Amaldeep & Sankaran, 2023). In this paper, they looked at these attacks and applied this to IPSec by comparing TLS/SSL and IPSec protocols. They then looked at attack vectors that could be employed in a similar way.

ALPACA – this attack exploits TLS servers implementing different protocols but using compatible certificates. Attackers can redirect traffic between subdomains giving valid TLS sessions without proper authentication[4].

DROWN – this is another TLS attack that enables threat actors to decrypt intercepted messages. It exploits the behaviour of SSLv2 servers using crafted (chosen ciphertext) messages to eventually get the victim's private key[5].

The area they had found to be most vulnerable was the IKE protocol (Internet Key Exchange) with IKEv1 being susceptible to DoS attacks (P. Dewan, 2008)(CVE-2016-5361). This is also relevant when looking at newer technologies like Cisco's AnyConnect (CVE-2023-20042) where an SSL/TLS session error allows for DoS conditions.

2.4 Configuration

In Khantamonthon and Chimmanee's paper, they covered the attack on Colonial Pipeline, which serves to further emphasise issues around configuration (Khantamonthon & Chimmanee, 2022). The attack involved a compromised password that had been used in in one of the VPN accounts. Not only was there no Multi-Factor Authentication, but the account in question had sufficient privileges that the hacker collective (DarkSide) was able to access the entire network (an issue with defence in depth that allowed for lateral movement) and deploy ransomware. This points to issues in applying access policies (i.e. least privilege).

Japanese game-maker Capcom suffered a breach of their VPN server (Khantamonthon & Chimmanee, 2022) which, given the attackers ability to move laterally throughout the network, resulted in over a terabyte of sensitive data being stolen. The initial breach was due to Capcom's continued use of a legacy device as their backup VPN server, which they were due to upgrade. This case is very relevant and highlights that breaches are not necessarily the result of recent CVEs or new attack vectors. There were clear oversights with legacy systems and, possibly more importantly, insufficient Defence in Depth principles had been applied. This allowed the attackers to move throughout the US and Japanese networks.

2.5 Client-side Vulnerabilities

[4] "ALPACA Attack," - https://alpaca-attack.com

[5] "The DROWN Attack," - https://drownattack.com

While enterprises are not definitely able to prevent, or pre-empt, client-side software attacks, awareness of the scale of the issue and applying patches in a timely manner is essential. This section also introduces the idea of mitigation through defence in depth, or more specifically through the application of Zero Trust Network Access as part of a DiD effort.

In a recent security report[6], ThreatLabz polled over 600 security/network engineers, found that 56% reported at least one VPN related attack in the previous year and stressed the escalation in VPN attacks. The focus of the report was on a need for a move to Zero Trust Networks given the severity of breaches involving VPNs.

CVE's:

CVE	Date	Attack
2023-21887 2023-46805	Oct 2023	Ivanti – issues with authentication check allowing for command injections
2024-22394	Jan 2024	SonicWall – one of numerous vulnerabilities reported in 2024 allowing for authentication bypass
2024-23112	Jan 2024	Fortinet – Authorisation bypass
2024-21888 2024-21893	Jan 2024	Ivanti – US federal agencies were ordered to discontinue use after CISA directive.
2024-23112 2023-4877 2023-47534 2023-36554 2024-42789	Mar 2024	Fortinet – multiple vulnerabilities reported that apparently allowed for remote code execution and crafted HTTP requests to be sent
2024-3400	Apr 2024	PAN-OS execution of code with root privileges on firewall

These attacks include some of the biggest names in network security, such as Palo Alto Networks and Fortinet and, while not mentioned, Cisco also had vulnerabilities in its AnyConnect VPN client in 2023 (CVE-20275, CVE-20241, CVE-20240, CVE-20178, CVE-20042). Other noteworthy attacks included CNA Financial, who were breached due to a compromised user account which led to lateral movement, privilege escalation, and eventually control of the network including devices connected by VPN.

2.6 Client-side Vulnerabilities

A prominent issue with VPN technologies is fingerprinting where techniques, such as a Counting Blook Filter with Chained Hash Tables (Wu, et al., 2022) or Machine Learning algorithms (Almutairi, et al., 2024) (Wang, et al., 2022), are used to identify traffic. This fingerprinting could be chained with other attacks to target organisations with sensitive data.

2.7 Future of VPN security

Given issues around encryption strength, particularly in light of quantum computing's arrival, and the difficulty in securing remote access devices, new applications for existing technologies may provide answers. Cisco has pushed for a move away from passwords to biometric authentication, with others repurposing (FIDO2) security keys to allow for their use with legacy systems (Huseynov, 2022). Quantum Key Distribution attempts to modernise VPN encryption

[6] https://zerotrust.cio.com/wp-content/uploads/sites/64/2024/05/threatlabz-vpn-risk-report-2024.pdf

by using QKD derived keys in the VPN setup (Buruaga, et al., 2023). While these are promising technologies, integrated approaches such as ZTNA seems to be the most sensible approach, with major vendors (such as Palo Alto[7]) suggesting it as the most realistic means of minimising risk. Defence in Depth will always be central in discussing VPNs.

3 Methodology

The goal of this project was to create a framework that outlined the most common issues leading to vulnerabilities when setting up VPNs in an enterprise environment. I took and inductive approach, as it would not have made sense to create categories for attacks and then look for research to support my initial hypothesis.

I could have taken a purely quantitative approach, using stats from the CVE database and other online resources, and then analysing this data. This would have meant severely restricting the scope of what I was hoping to achieve and, realistically, reducing my engagement with the technology involved. It would have meant that I focussed only on exploits in the hardware/software, and not in the configuration and application of VPN technologies in the enterprise setting.

Alternatively, I could take a qualitative approach and look to build labs to analyse various attack vectors. This posed the same problem as that of a purely quantitative approach – namely, having to restrict the types of attacks and focus on a proportionately small number of issues given the sheer number of attacks to be covered. This would have afforded me huge exposure to the technology but would not have been representative of the threat landscape as a whole.

I chose a hybrid approach, using a combination of recorded cyber security incidents, research papers, recommendations from industry leaders (IBM, CISCO, NIST, etc), reports, literature based around the technology, and the CVE database. This gave more insight into the data and allowed a better understanding of vulnerability patterns in ways a purely quantitative approach couldn't. Or, to put it another way, it allowed to answer "how and why" questions when the data is pointing at a particular vulnerability (Penta & Tamburri, 2017). This mixed method has meant that the ability to parse data and correctly retain/remove data points is paramount in achieving accurate results. The criteria I used was:

- The attack needed to be the result of oversights on the part of the organisation (not the provider of the technology), except in the case of patching
- The attack vector must involve virtual private network technologies
- The attack must be verifiable – either via a CVE, research paper, K8 filing, or lab
- The attack must be relevant, i.e. involve systems that are still in use

3.1 Collecting data

Research papers provided the foundation for my understanding in the technologies and challenges faced. I then used attack data to create the categories for this framework. I limited the CVE search to attacks on VPNs from 2020 onwards, as attacks prior to this are still caught in the "legacy systems / patching" category.

[7] https://www.paloaltonetworks.com/resources/guides/zero-trust-overview

CVE's - This involved searching through the CVE database using keywords and searching by product and vendor (a full list of search terms and results is provided in Appendix). This gave a total of (556) results for VPN related search terms.

Research Papers - of over 100 papers reviewed, close to 50 were used in understanding this topic (again, list provided in Appendices).

Books - covering VPN technologies and encryption, this gave a greater understanding of the subject as well as outlining best security practices and common issues.

Labs - were used to verify attack mitigation strategies, though this was limited in scope as many attacks weren't feasible to recreate.

3.2 Clean and Sort the Data

This involved removing duplicates and any irrelevant data, which was done in excel (included in uploaded files), and then reviewing all data points.

The data was divided into client-side and enterprise side vulnerabilities. Then the attacks were sorted based on whether they were configuration issues (organisational responsibility) or product issues (vendor responsibility).

3.3. Analyse data

In collecting data, I started with the CVE database, though many breaches in the real world are the result of misconfiguration. In this regard, the CVE's tended to over-represent some aspects of attacks, while under-representing others. It also begged the question – whose responsibility is this? CVE-2024-21606 highlights an issue with the Juniper OS where a “double free” allowed for arbitrary code execution. There was nothing the enterprise could have done, besides timely updates.

The table below shows the recorded vulnerabilities for each year from 2020. A clear increase in client-side vulnerabilities indicates an increase in attacks, but it is only when combined with industry reports that we can start to appreciate the figures.

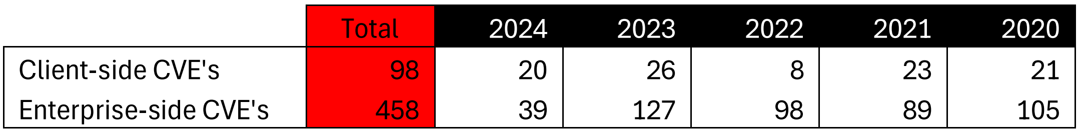

	Total	2024	2023	2022	2021	2020
Client-side CVE's	98	20	26	8	23	21
Enterprise-side CVE's	458	39	127	98	89	105

VPN Based CVEs from 2020

At this point, I used industry reports, RFCs to verify technical details, books, and guidelines to validate the data. Realistically, this allowed a degree of subjectivity into the results as, while all points are backed by research, there wasn’t sufficient attack data to enable statistical models to be applied.

There are also difficulties in recording issues with “unpatched” systems insofar as there is no CVE for this, so it is essentially based on empirical data, and this can only be classified via reports.

Taking these metrics together can help to paint a clearer picture of vulnerability and attack trends, that separately would not be possible to outline. CVE's track vulnerabilities “in the wild”, research proposes vulnerabilities and best practices for mitigation, and industry feedback shows trends and gives granular (anonymised) data on the frequency and types of attacks experienced.

4 Design Specification: Frameworks

To develop this framework, I needed to first understand what underpins the frameworks that are used industrywide. I chose the 4 frameworks below as being demonstrative of cybersecurity frameworks in general, as they give a good cross section of complexity, security areas, and technologies. Frameworks, ideally, should have the CIA triad at their core.

MITRE – Characterised by it's almost overwhelming breadth and depth, it provides security professionals a way of charting attack vectors in a very structured manner. There are 16 categorised steps. It is not applicable to most save SOC Analysts.

NIST CSF – Divided into 6 functions, each with its own categories and subcategories, this is a robust framework that is clear in its intentions. It provides a concise introduction (30 pages) and gives example implementations, as well as a GUI and references to other frameworks.

OWASP RAF/Top10 – These are probably the most interesting projects from a framework point of view. The Risk Assessment Framework (RAF) comprises a pair of testing tools with a user guide. The idea was to simplify the testing process by providing pre-built tools that replace combinations of tools that can be complex to setup. The Top10 is not a framework, but a standard awareness document.

ISO 27000 – Framework for information security management systems (ISMS). This is a substantial undertaking. There are up to 93 security controls and multiple guidance documents. Words such as assessing, establishing, implementing, and maintaining are often used in describing ISO standards.

CIA Triad – While technically not a security framework, it functions as the foundation of cybersecurity and illustrates the appeal of simplicity. It allows for flexibility but doesn't offer the structured approach that a procedure-oriented framework would.

While a lot of ground is covered by these frameworks, there are consistencies that stand out as being essential in making them work:

- Clear and documented policies / procedures
- A measurable reduction in risk (solving a specific problem)
- Viable solution (one that is applicable for / available to any enterprise with a VPN)

In this sense, the starting point for a model would include:

1. Introduction outlining the goals
2. Assessment of current situation
3. Analysis of issues recorded
4. Outline mitigation steps
5. Re-assess, Review, and document

5 Implementation

5.1 – VPN DOS Attack

The lab shows a vulnerable edge router that has been exploited by creating an adjacency. This allowed for route poisoning, which meant the gateway was accessible. At this point:

- The attacking machine can send crafted ISAKMP packets to cause a gateway DoS
- The edge router must be made “passive” to protect from route poisoning
- The gateway should be upgraded to IKEv2 which is not vulnerable to the DoS attack

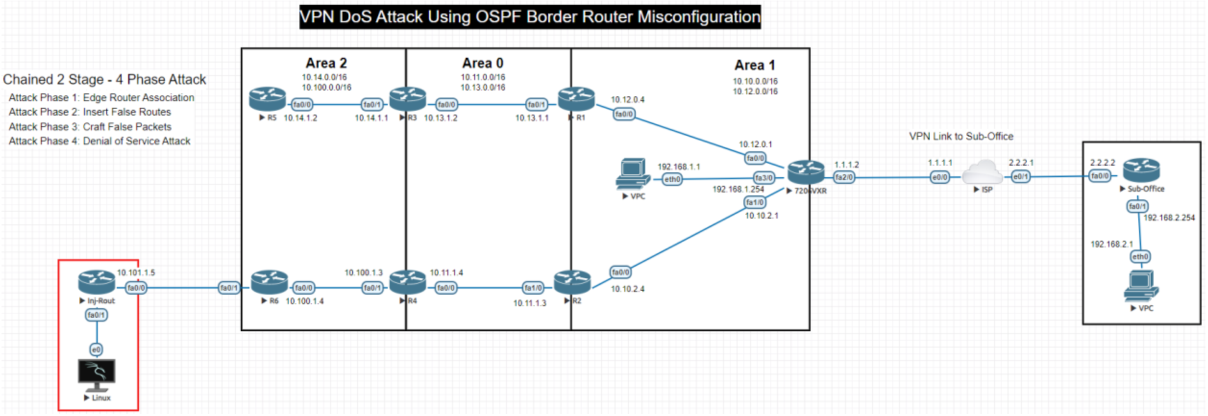

5.2 - The Framework in Practice

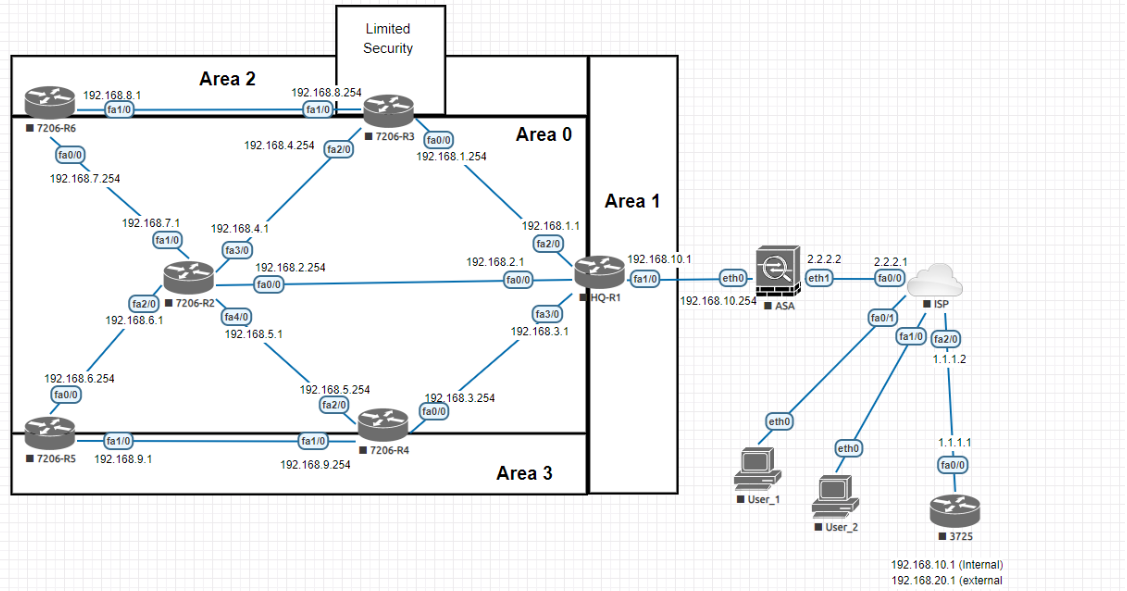

This is a similar, but more realistic scenario depicting a conventional enterprise network. The assessment is done to outline all vulnerabilities and then analysis allows for risk scores to be applied. The process is covered in more detail in the Evaluation section, but:

- The lab covers a wide range of issues that had been highlighted in the research
- After finishing the assessment, attack simulations should be used to verify success

6 VPN Framework and Evaluation

Applying this framework in the lab scenario provided a clear and measurable reduction in risk. This was demonstrated using Wireshark, Scapy, and configuration updates - best practices as advised by industry leaders, and then attempting attacks. The labs were designed to mimic real world networks and are particularly useful for applications in government organisations, where networks will have grown organically over large periods of time.

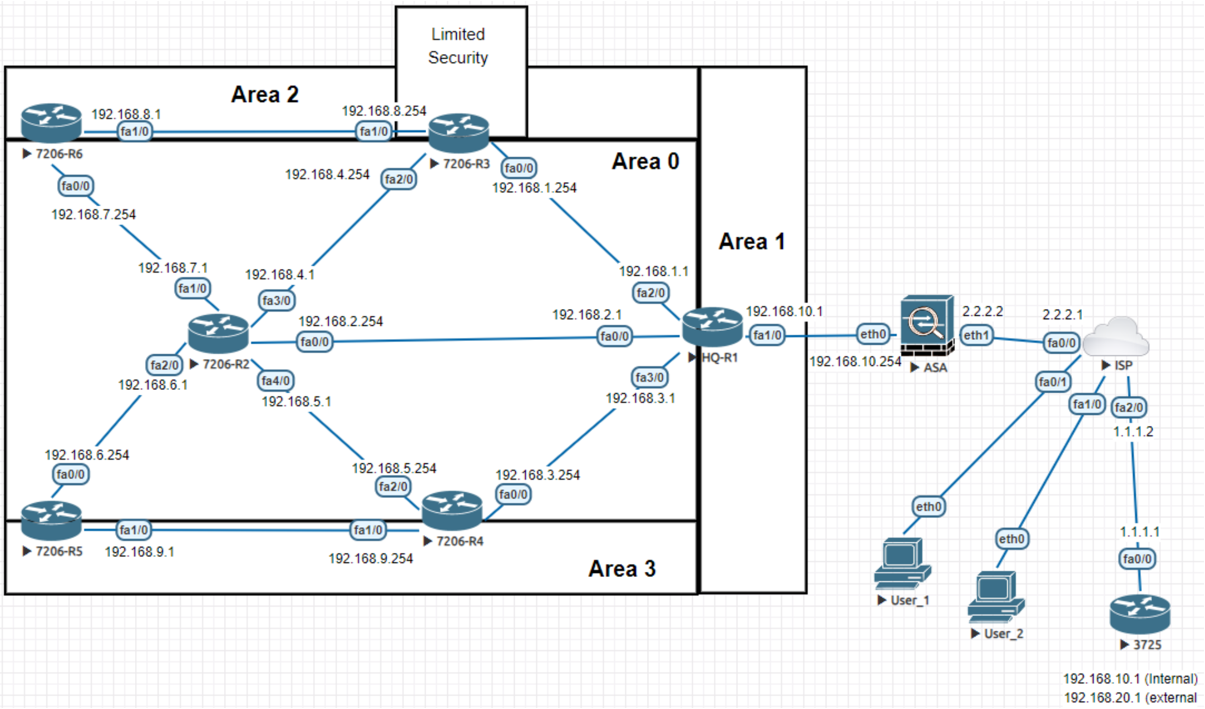

Lab 2 – Framework Model

In the first lab scenario, we see how route poisoning and a DoS attack could be perpetrated without proper configuration guidelines. This was to validate the core concepts and demonstrate assessment - mitigation techniques in the lab environment. In the second lab (shown above), we apply the framework and look at mitigation techniques. In outlining the protocols and processes in place for each of the categories, we can quickly spot high-risk areas. The risk scoring system is easily applied but limited in scope, with 3 main risk brackets used (more details provided in the configuration manual). The assessment involves assigning risk scores to each of the 18 assessment areas. The process of applying the framework is shown in the diagram below, where we look to clarify the company’s appetite for risk and how this is reflected in their Business Continuity or Disaster Recovery Plan. Then the assessment is done, and analysis undertaken which allows for risk scoring. Mitigation steps are outlined and then scheduled, with reviews to allow for re-evaluation of risks.

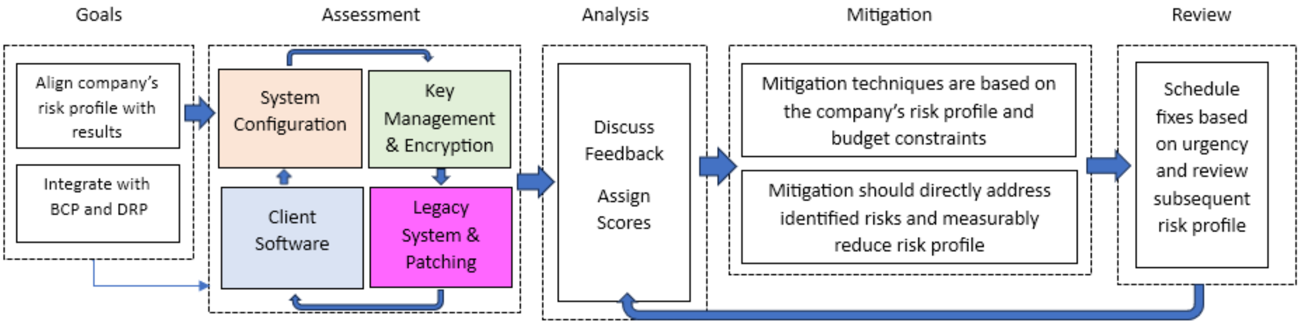
Goals
Assessment
Analysis
Mitigation
Review
Align company's risk profile with results
Integrate with BCP and DRP
System Configuration
Key Management & Encryption
Client Software
Legacy System & Patching
Discuss Feedback
Assign Scores
Mitigation techniques are based on the company's risk profile and budget constraints
Mitigation should directly address identified risks and measurably reduce risk profile
Schedule fixes based on urgency and review subsequent risk profile

Framework Process Flowchart

The four categories assessed in this framework are:

System Configuration	Key Management and Encryption	Legacy Systems & Patching	Client/3rd Party Software
VPN Protocol Security 1 VPN Protocol Security 2 Default Config (Access) Default Config (Traffic) Firewall Rules Unsecured Access to Devices	Tunnelling Key Management Encryption Standards	Software Checks Patching Schedule Inventory Monitoring Legacy Systems	Password Management Access Policies (Network) Access Policies (Devices) 3rd Party Access Network Segregation

Framework Categories Table

By applying the framework in the lab, numerous vulnerabilities were uncovered, and mitigation strategies were drawn up. In the 4 areas above, risk scores of 27 (SC), 14 (KME), 21 (LSP), and 23 (CS) were given. This meant a score of 85 out of a possible 170. Key risk areas were:

- Password policies on core systems, including a pre-shared key for the VPN
- Physical access to equipment (remote sites not physically secured)
- Old Key Management protocol (IKEv1)
- Weakened encryption standards (due to outdated key management)
- No patching schedule and legacy devices still active and unmonitored
- 3rd party VPN access does not require MFA

There are some serious issues here, and an initial risk score of 85 (of a possible 170) was given. This was then coupled with CVE numbers to stress the severity of some of the issues.

A CVE score is included (the number of CVE attacks showing for that given product or process) which helps to identify places that provide a large threat surface. Details of the scoring, and CVEs are provided in the uploaded data_file.xls [Framework_Sample tab].

	Before	After
Risk Score	85/170	37/170
CVE Score	64	0

Risk Scores Pre- and Post-Mitigation

6.1 Mitigation steps

By using the scoring system to quantify the severity of the risks, it becomes possible to prioritise high-risk areas. In this way, we can outline an overarching strategy and allocate resources to high impact areas, maximising return on work hours and financial costs. Each assessment point with a score greater than 5 should have a plan for remediation.

6.2 Framework Issues

The results of this framework while promising, would still not be sufficient for industry application.

6.3 Application of the Framework

As shown above, the scoring is based on self-assessment, which is one area of this framework that would need reviewing. A more meticulous system for scoring would provide greater clarity but may also prove overwhelming.

Recommendations – while every organisation would benefit from the application of this framework, the number of possible scenarios to be accounted for makes it unrealistic to provide discrete mitigation recommendations in all cases. This framework aimed to provide best practices relating to VPN standards used and potential configuration pitfalls, but the scope has meant there will be scenarios that aren't covered.

6.4 Data Gathering & Replicability

One of the difficulties with building this framework was in modelling the data. It was possible to categorise and analyse, but in trying to apply statistical models – even something simple such as a linear regression, it lost meaning. Trying to predict attack vectors based on changes in an independent variable, or even quantifying future attacks becomes very difficult without adopting a more rounded approach. This lends a degree of subjectivity to any results, as they are always interpreted through the lens of the research and data chosen. In this regard, replicability could be difficult as different research and technologies may be favoured, resulting in different recommendations.

With more data points, and more time, I think modelling may show relationships in attacks and be useful in predictions. Realistically, though, this may be an opportunity for AI to shine.

6.5 Discussion

There are several areas that warrant discussion given the goals of this framework. The recent explosion of all things AI does present opportunities, and possibly risks, for VPN implementations. The idea of AI VPNs has already been broached by Zscaler and it does seem promising in reducing the difficulty and overhead in deploying VPNs. This is especially true with initiatives such as Civilsphere's AI-VPN[8] that aims to help activists and journalists protect their traffic. That said, there are still concerns around data privacy and the fact we will never be how the implementation is being managed "under the hood".

Discussions around Open-Source software has also recently taken a front seat given the proposal of new European technology laws. The implications of this for commonly used open-source VPN products like Open-VPN will probably be negligible, but there may be other unintended consequences.

[8] https://www.civilsphereproject.org/aivpn

This framework was developed primarily with the current threat landscape in mind (outside of known issues in encryption etc), and this means the lifespan of the framework may be limited.

There may also be issues with the scope of the framework, in that there is always a substantial amount of crossover when dealing with multi-faceted technologies such as VPNs, and knowing where to draw the line can be difficult. This would have implications when incorporating this project into an overarching security framework in that there may be overlap, or worse, gaps in pairing it with neighbouring technologies such as automating networks.

The Framework itself was well researched and developed but was somewhat lacking in concrete recommendations when dealing with threats. It needs a clearer system of quantifying risks, allowing those applying the framework to better understand the risks and to better prepare for mitigation tactics.

7 Conclusion

The goal of this project was to create a framework that could significantly reduce the risk of, and hopefully impact of, successful attacks on VPN implementations. Realistically, a lot more work would need to be done in terms of data gathering and building out the assessment, but the kernel of something very useful is still there.

As we move towards Zero Trust models, VPNs will still be necessary and possibly central to how we work. With that in mind, a VPN framework (such as this), with Defence in Depth as a core component, will prove invaluable in limiting the severity of unavoidable security flaws when using VPN technologies.

References

Abbas, H. et al., 2023. Security Assessment and Evaluation of VPNs: A Comprehensive Survey. *ACM Computing Surveys.*

Almutairi, S., Neumann, Y. & Harfoush, K., 2024. Fingerprinting VPNs with Custom Router Firmware: A New Censorship Threat Model. *2024 IEEE 21st Consumer Communications & Networking Conference (CCNC),* pp. 976-981.

Amaldeep, S. & Sankaran, S., 2023. Cross Protocol Attack on IPSec-based VPN. *11th International Symposium on Digital Forensics and Security (ISDFS),* pp. 1-6.

Buruaga, J. S. et al., 2023. VPN Protection with QKD-Derived Keys Using Standard Interfaces. *2023 23rd International Conference on Transparent Optical Networks (ICTON),* pp. 1-4.

Huseynov, E., 2022. Passwordless VPN using FIDO2 Security Keys: Modern authentication security for legacy VPN systems. *2022 4th International Conference on Data Intelligence and Security (ICDIS),* pp. 01-03.

Jahan, S., Rahman, M. S. & Saha, S., 2017. Application specific tunneling protocol selection for Virtual Private Networks. *2017 International Conference on Networking, Systems and Security (NSysS),* pp. 39-44.

Jahan, S., Rahman, S. & Saha, S., 2017. Application specific tunneling protocol selection for Virtual Private Networks. *International Conference on Networking, Systems and Security (NSysS),* pp. 39-44.

Khantamonthon, N. & Chimmanee, K., 2022. Digital Forensic Analysis of Ransomware Attacks on Virtual Private Networks: A Case Study in Factories. *6th International Conference on Information Technology (InCIT),* pp. 410-415.

Mahmmod, K. F., Azeez, M. M. & Ahmed, M. A., 2020. IPsec Cryptography for Data Packets Security within VPN Tunneling Networks Communications. *International Conference on Electrical Engineering and Informatics (ICELTICs),* pp. 1-8.

P. Dewan, D. K. S. N. a. D. M. D., 2008. Denial of Service Attacks Using Internet Key Exchange Protocol. *5th IEEE Consumer Communications and Networking Conference,* pp. 471-475.

Pedapudi, S. M. & Vadlamani, N., 2022. A Comprehensive Network Security Management in Virtual Private Network Environment. *2022 International Conference on Applied Artificial Intelligence and Computing (ICAAIC),* pp. pp. 1362-1367.

Penta, M. D. & Tamburri, D. A., 2017. Combining Quantitative and Qualitative Studies in Empirical Software Engineering Research. *IEEE/ACM 39th International Conference on Software Engineering Companion (ICSE-C),* pp. 499-500.

Rao, M. S. et al., 2023. The Use of Virtual Private Networks to Facilitate Remote Working: A Critical Review of Cyber Security Implications to Develop Successful VPN Systems. *2023 3rd International Conference on Advance Computing and Innovative Technologies in Engineering (ICACITE),* pp. 592-595.

Sawalmeh, H., Malayshi, M., S., A. & Awad, A., 2021. VPN Remote Access OSPF-based VPN Security Vulnerabilities and Counter Measurements. *International Conference on Innovation and Intelligence for Informatics, Computing, and Technologies (3ICT),* pp. 236-241.

Snader, J. C., 2006. *VPNs Illustrated: Tunnels, VPNs, and IPSec.* s.l.:Pearson Education.

Tay, W. J., Lew, S. L. & Ooi, S. Y., 2022. Remote Access VPN using Mikrotik Router. *2022 International Conference on Computer and Drone Applications (IConDA),* pp. 119-124.

Tongkaw, A., 2021. VPN Security in Campus Network during Covid-19 epidemic: Case Study in Southeast Asia. *International Conference on Electrical, Computer and Energy Technologies (ICECET),* pp. 1-6.

Uskov, A. V., 2012. Information Security of IPsec-based Mobile VPN: Authentication and Encryption Algorithms Performance. *IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications,* pp. 1042-1048.

Wang, Y., Yu, G., Shen, W. & Sun, L., 2022. Deep learning based on byte sample entropy for VPN encrypted traffic identification. *2022 5th International Conference on Advanced Electronic Materials, Computers and Software Engineering (AEMCSE),* pp. 293-296.

Wu, H., Liu, Y., Cheng, G. & Hu, X., 2022. Real-time Identification of VPN Traffic based on Counting Bloom Filter and Chained Hash Table from Sampled Data in High-speed Networks. *ICC 2022 - IEEE International Conference on Communications,* pp. 5070-5075.

Xu, Z. & Ni, J., 2020. Research on network security of VPN technology. *2020 International Conference on Information Science and Education (ICISE-IE),* pp. 539-542.