WHITE PAPER:

Remote Access—Attack Vectors

Threats, Findings & Remedies

NCP
SECURE COMMUNICATIONS

Table of Contents

Overview

The variety of mobile devices on company networks has eased communications across long distances, but has also increased the risks. Organizations sharing data and collaborating on inexpensive and unverified infrastructures face evolving and sophisticated attackers. While virtual private networks (VPN) and secure remote access are not new, the threats they combat are constantly changing and require regular monitoring and security updates to stop.

This paper provides a fresh look at the threat landscape and offers several approaches for keeping data safe in the constantly moving world of remote connectivity. It outlines examples of real-life attacks that successfully eschewed security protocols and then reviews some of the latest VPN vulnerabilities, including flaws and root causes. It also provides guidance on configuring and managing VPNs to reduce the chance of a breach.

Threats to Remote Connectivity

A VPN provides protection, but it also presents an attractive target to attackers for two reasons:

First, the VPN transmits sensitive information over public and shared networks. The extension of the network outside the perimeter makes assets much more accessible for attackers. The hackers no longer have to breach a perimeter to find data that is sensitive. Just one remote access flaw could be sufficient to breach a system. Organizations with internal networks that store highly sensitive data increase their risk of exposure when they use VPNs to communicate. Hackers are drawn to the vulnerability VPNs present when private information is transmitted and, thus, made more accessible.

Second, a VPN typically does not have layers of security found in perimeter defenses, yet it will provide access from outside a perimeter to inside networks. Firewalls, intrusion detection, proxies and other controls create a barrier for traffic that flows from public networks to private and even a barrier for outbound traffic. A VPN often will lack many of these controls; an internal system might communicate over a private network to a hostile system on the outside. Attacks from a public network will then succeed when they would have been prevented or at least detected. An internal network often has systems that are less secure than those expected to handle traffic from the public network. This can make VPN-based attacks that bypass a perimeter more attractive than attacks that directly target the perimeter.

Due to these two threat factors, VPN systems have to address numerous vulnerabilities. Attackers constantly try to find ways to exploit hardware and software, often taking advantage of misconfigurations and poorly managed implementations. This is why VPNs continue to be involved in serious breaches.

Breach Examples

The most infamous example of a VPN-related breach was the 2008 incident at Heartland Payment Systems. The payment processing system, which obviously contained highly sensitive data, was breached, in part, using a VPN. First, an SQL injection attack allowed hackers to take control of a system inside Heartland, but not within the transaction processing environment.[1] The attackers then used the internal system to elevate their access. They were able to do this because systems, such as workstations, were not secured to the same level as systems that hosted sensitive data, but they were only a step away. The less secure systems connected to the transaction processing environment were compromised, and sensitive information was stolen.

A similar incident affected Google in January 2010.[2] Following the incident, Google instructed its users to make emergency changes to VPN settings due to highly sophisticated and targeted attacks suspected to be from China.[3] The attacks later were said to be directed at un-patched browser flaws.[4] The timing of the VPN configuration change led some to conclude there was a relationship between internal systems with un-patched software and VPN access to internal networks and sensitive data.[5] Segmentation, proxies and filtering, as well as monitoring, may have prevented this breach. Patching (upgrading the browser) also may have prevented it. The suggested changes show how layers of security and good management—and not just settings—are critical to the security of a VPN. Ironically, three months after the attack, Google suggested that its users in China consider using VPN technology to bypass perimeter security and to ensure ongoing access to Google services in mainland China.[6]

Another example of insufficient VPN management and security that lead to a breach comes from an employee terminated by a utility company, Energy Future Holdings. The employee was able to use the VPN—even after his position was terminated—to access the corporate systems used for consumer demand forecasts.[7] The terminated employee used the access to corrupt data, which caused $26,000 in lost business alone. This breach reinforces that it is absolutely essential to cease VPN access for terminated and former employees as quickly as possible. The best way to do this is connect user provisioning and identity systems with VPN administration.

1 http://www.wired.com/images_blogs/threatlevel/2009/08/gonzalez.pdf

2 http://chenxiwang.wordpress.com/2010/01/21/ok-there-is-more-or-may-be-less-to-the-vpn-story-google-says/

3 http://www.sophos.com/security/topic/operation-aurora.html

4 http://www.microsoft.com/technet/security/advisory/979352.mspx

5 http://googleblog.blogspot.com/2010/01/new-approach-to-china.html

6 http://knol.google.com/k/circumventing-the-great-firewall-of-china#Tools_for_Circumventing_the_GFW

7 http://www.wired.com/threatlevel/2009/05/efh/

Retail organizations should have extensive experience managing access to systems, given the high turnover of employees in the retail industry. Even so, Walmart, one of the largest retail chains in the world, left a VPN connection open to an employee that no longer worked for the retailer in 2005.[8] Unlike the previous breach that lasted just one day, the Walmart violation ran for 17 months. During this time, the attacker accessed payment card systems. The breach was only detected when a failed installation of a password-dump tool caused a system to go offline. The outage alerted operations staff for the first time. It took a breach of this significance for the attacker's VPN account to finally be closed. But the attacker was still able to reenter the system. Again, the best way to combat such a crisis is to connect user provisioning and identity systems with VPN administration. An additional fix in this case is segmentation, which adds additional controls required for VPN access to sensitive data.

The lesson from the Walmart breach should have prevented other retailers from making the same mistake. Yet three years later in 2008, a similar incident happened at a medical device manufacturer, Stryker Instruments.[9] An attacker, suspected to be a terminated employee, connected with a VPN multiple times over several months. When the organization discovered the breach, its customers (across 48 states) had to be notified. The company's official notification letter pointed to flaws in the management of VPN, as well as the need to move to a more secure VPN.

While Stryker suspected a terminated employee, the organization was not able to identify the culprit for certain. This could be because of improper disposal of VPN equipment, which was the case with Kirklees Council.[10] In this scenario, an IT security professional purchased a used VPN Gateway for less than $2. After setting up the device, he found that it was still configured to access the prior owner's internal network.

Causes

A close look at vulnerabilities involved in breaches, like the examples described above, can be divided into three categories for analysis.

Quality: VPN systems are expected to handle complex security operations in an un-trusted environment, especially if they are offered as an advanced form of network privacy. Flaws can be expected in this situation unless particular care and testing is taken during engineering. That means not all VPN products are created equal, and they may be distinguished by their emphasis on quality and security.

Design: Second, the level of security varies by customer. Some require more strict control than others. Therefore, default configurations might emphasize ease of deployment and connectivity more

8 http://www.wired.com/threatlevel/2009/10/walmart-hack/

9 http://www.wired.com/threatlevel/2009/10/walmart-hack/

10 http://doj.nh.gov/consumer/pdf/stryker_instruments.pdf

than security. The aim of some VPN products is to extend private networks across distance for reduced cost compared to physical lines, with security as an option. Not all VPN products are setup with the same defaults or allow the same options.

Management: The complexity of networks typically requires specialized and professional deployment, configuration and management. A VPN is part of the network, but it also calls for security domain expertise with authentication and authorization systems, coupled with protocols, such as SSL and IPSec, that few are compelled to learn well. Differences between Diffie Hellman[11] Group 1 and Group 2, for example, will be unfamiliar to most although choosing the wrong one will show up on a simple vulnerability scan of an IPsec VPN. Another example is a VPN device accidentally left enabled or an account authorized after an employee has been terminated. Finally, a system that runs vulnerable software can be exploited while connected to a VPN, giving attackers access to private networks. VPN deployments can be deceptively complex and must be managed well to provide appropriate levels of security.

Although the vast majority of breaches involve management issues, design and quality are still very important considerations. Both design and quality are among the best ways to differentiate VPN products and solutions. The following section details several vulnerabilities in quality and design.

Vulnerability Examples

Vulnerabilities in VPN can often be found on vendor support sites as well as security notification bulletins such as Secunia[12] and the U.S. NVD.[13] Vendors will notify customers of the vulnerability after a fix has been created, but advance notice is sometimes possible by monitoring disclosure and security forums.[14] Advance notice is often useful when vulnerabilities are urgent or critical in severity and will require planning or immediate attention.

Information Leaks

After the installation of VPN software, network traffic might still exit through other interfaces. A high-profile vulnerability was announced in 2010, caused by the combination of IPv6 and PPTP-based VPN, and exposed users' IP address, MAC address and computer name.[15] The Findnot.com VPN in 2006 was found to redirect traffic unencrypted to the public net when its servers were congested,

11 http://news.bbc.co.uk/2/hi/technology/7635622.stm

12 http://www.rfc-archive.org/getrfc.php?rfc=2409

13 http://secunia.com/community/

14 http://nvd.nist.gov/

15 http://seclists.org/fulldisclosure/; http://oss-security.openwall.org/wiki/mailing-lists/vendor-sec

without notifying its users. A similar issue may occur after the installation of multiple VPNs that do not work well together and therefore fail to isolate traffic, causing information leakage across the VPNs. This was the case with Cisco VPNs that had a bug when processing extended communities.[16] Cisco's VPNs incorrectly used a corrupted route target (RT) to forward traffic, causing a leak—from one VPN to another. The solution to these leaks, aside from vendor patches, is hardening a system specifically to prevent route failures, DNS leaks and IP leaks outside a VPN connection, even after the preferred VPN connection fails.

Another recent vulnerability in this area is related to clientless VPN products. These devices retrieve content from different sites and then serve the data so it appears to be from the SSL VPN. This circumvents same-origin restrictions. A same-origin policy is designed to enforce trust domains and block malicious scripts. The SSL VPN design breaks this by trying to present remote site data with different domains as its own. An attacker with a malicious page viewed through a clientless VPN thus can hijack a user's session or capture keystrokes. The solution to this is to configure VPN clients to access only trusted domains.

IKE aggressive mode (AM) has a vulnerability that can lead to a serious leak of information. Attackers are able to extract password hashes and attempt password brute force methods due to different responses returned for valid and invalid usernames. This is prevented by configuring devices to deny and never initiate any IKE AM connections.

Caching and Duplication

VPN client programs often store authentication credentials (e.g. username and password) to make network access more convenient. This may even be a default setting in some VPN products, making sensitive data extremely vulnerable. An attacker may search for cached credentials to exploit a VPN, which can be as simple as looking for plain-text passwords in memory or the registry. A 2007 vulnerability with the CheckPoint VPN-1 SecuRemote/SecureClient AutoLogin feature, for example, had to do with authentication credentials stored in the Windows registry (subkey 'Credentials' in HKLM\Software\Checkpoint\SecuRemote). A local user could easily access the credentials to authenticate to a VPN as a target user.[17] A similar problem is, some clients use non-encrypted (obfuscated) passwords, which are very easy for attackers to find. Likewise, sensitive data may be left behind on a system after a VPN session ends. Passwords should be protected properly at all times, and the cache of a client that has accessed a VPN should be cleaned after every session.

16 http://www.wired.co.uk/news/archive/2010-06/18/huge-privacy-flaw-found-in-vpn-systems

Duplication or replay of VPN traffic also continues to be an important vulnerability. Data may be encrypted in communication and signatures used to authenticate every packet, but a motivated attacker may still perform a replay with success. VPN implementations of IPsec apply a sequence number to packets, using the Encapsulating Security Payload (ESP). A table of these are maintained by the destination host and then used to monitor for duplicate packets and invalid traffic. This system works only when the sequence number is also encrypted or signed. Otherwise, it is possible for an attacker to monitor and figure out the sequence patterns.

Both of these vulnerabilities demonstrate the importance of quality in VPN products as well as the need for careful design.

Denial of Service

Packets that are malformed or those that do not conform to RFC specifications, can cause overflows in buffers and terminate services by exhausting resources. This affects a VPN like any other network system.

SSL-based VPNs, for example, may be based on a standard TCP connection. They can be configured to use UDP instead, perhaps to get higher performance. The use of UDP, however, makes blocking and detecting denial of service attacks much more difficult. UDP also makes it very easy to detect a VPN on a network (TCP-based traffic is more easily hidden—traffic on port 443 looks like HTTPS). One attack on Nortel VPNs was due to a malformed ISAKMP header that was undetectable by IDS because it was so similar to a standard IKE header.[18] The hidden attack packets would cause an immediate crash that would either force the device to reboot (five minute outage) or cause it to hang indefinitely. Another view of denial of service is resource exhaustion by a flow of many packets that cause a VPN server to run out of negotiation slots. This makes a server unable to process legitimate connections. The Cisco VPN 3000 was found to be vulnerable if bad traffic was sent on a particular port, in this case, port 80. However, only a short stream of less than 50 packets was needed to hang the device.[19]

Remote Connectivity Security Options

Anyone who uses a VPN in a shared network environment for wide-area connections, especially with mission-critical applications, will benefit from a set of basic security requirements. These requirements will apply regardless of the technology used to create a VPN.

17 http://www.cisco.com/warp/public/707/cisco-sa-20080924-vpn.shtml

18 http://supportcenter.checkpoint.com/supportcenter/PublicLoginRedirect.jsp?toURL=eventSubmit_doGoviewsolutiondetails=%26solutionid=sk34315

19 http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-1802

Authentication

VPN client authentication by the VPN server is critical to effective security. Authentication typically happens at two levels, a certificate or a secret. Certificate authentication is done by exchanging certificates and/or pre-shared keys for the VPN endpoints. Secret, also known as password, authentication can be a database, or it can be tied to authorization in a central directory, such as Microsoft Active Directory, RADIUS or LDAP.

Two-factor authentication is required in environments that require high levels of security. The combination of something a user knows (secret/password) and an authentication certificate or token will stop many VPN attacks that are based on compromised passwords. Many VPNs support a token that generates a one-time password, which adds specific hardware as a requirement to be authenticated properly. Strong message digest algorithms, such as SHA, are also required to protect the validity of user authentication information, discussed below in the Integrity section.

Authorization

A user or system is allowed to access certain network resources and use the VPN only after they successfully pass authentication. Common authorization systems for integration are Microsoft Active Directory and RADIUS. The use of a centralized system cannot be understated, as it enables better VPN management and access. Conversely, authorization must be regularly reviewed and removed as soon as possible to protect a system from attack as described in the above breach examples.

Auditing

VPN logs are useful not only for troubleshooting, but also for detecting and responding to incidents. A list like the following can be used to configure the specific events to be collected, reviewed and archived:

1. User
2. Date, time and command
3. System location
4. Authentication success/failure
5. Authorization success/failure
6. Configuration change, especially to protection (anti-virus and intrusion detection)
7. Privileged access
8. Network addresses and protocols

An investigation of a VPN attack will depend on audit trails, since the details for each of these areas are essential to establish who did what, where and when.

Availability

Availability refers to the percentage of time that a VPN service is working. Connections that become unavailable run the risk of information leakage or denial of service. It is best for a VPN system to monitor its own health and migrate connections to a new VPN service without requiring administrator intervention. The failover also should allow users to continue their work without data corruption or the need to authenticate again.

Integrity

Integrity is handled by controls that prevent data modification. VPNs use three methods to provide integrity controls.

The first is a one-way hash function, which can take any length of input and create a fixed-length output value. This value is easy to calculate for a file, but it is difficult to create a file that would give the same value. Data, therefore, becomes difficult to modify without detection. The recipient of a file can validate its integrity by creating a hash, and comparing it to the hash that was created by the sender and sent with the file. A VPN can provide integrated hash mechanisms to generate a unique signature for the data in transit—when data is modified its signature will not match that of the unmodified data. Two algorithms frequently used by VPNs to verify data are MD5 and SHA. MD5 has been proven to be unreliable so the latter is the best option.

The second method is with message-authentication codes (MACs) that add a key to a hash functions. The MAC is calculated with a key shared between sender and recipient. The sender generates the MAC from a file and sends it with the file to a recipient.

Finally, digital signatures also can provide VPN data integrity. The digital signatures work in reverse compared with public key cryptography. A document is digitally “signed” by the sender using their private key, and then the recipient uses the sender’s public key to verify the document.

Confidentiality

Data is encrypted by a VPN to prevent disclosure when it is transmitted over a shared or public transit network. Key management and strong encryption are the foundation to ensuring encryption is effective. Length of an encryption key, for example, is an important decision. The largest reasonable size is recommended. .

Rethinking Remote Access with NCP engineering

NCP engineering provides a software VPN platform solution that is designed for an organization that requires control over large networks. The technology enables true network transparency, control and monitoring of VPN connections to and from a central point in the network. It supports both IPsec and SSL, and is the endpoint for secure communication with all mobile and stationary devices as well as remote gateways and branch offices.

The most comprehensive, easy-to-manage VPN solution in the North American market, it was built from the ground up. Seamless interoperability with existing infrastructure is available. The modular solution combines both SSL and IPSec management with strong policy enforcement. End-users can connect from any device to the network through a secure VPN tunnel with one click, regardless of if they are on a WLAN, LAN or cellular network.

It integrates fully with Cisco, Juniper, Check Point, SonicWall and other major vendor equipment, preserving existing technology investments for customers.[20] Practical NAC function support and management is also provided with the installation of the complete enterprise system.

NCP Secure Enterprise Management System

A centrally-controlled software solution that provides network administrators with a single point of administration for a company's entire IPSec and SSL VPN network, as well as practical NAC functions. All status information is made graphically available on the system monitor in real time, and plug-in updates and configuration settings can be easily controlled and distributed. User data can be imported via standardized interfaces from existing directory services and identity and access management systems (IAM). Built-in transition software ensures redundancy systems guarantee high availability of the management system, avoiding costly downtime and loss of policy settings.

NCP Secure Enterprise Server

A hybrid IPSec and SSL gateway that controls and monitors all VPN connections to and from the central data network. It offers high availability clustering to maintain network performance speeds and allows administrators to run up to 10,000 concurrent VPN sessions. Unique to NCP, the gateway provides one plug-in for full network access. The NCP Secure Enterprise Server supports the industry's widest variety of endpoint platforms and any IPSec-based device, including the iPhone.

NCP Secure Enterprise Client

A bundled client, personal firewall and connector provide the most secure end-point connection for the industry's widest array of platforms, including Windows-based (Mobile 5/6x, CE, XP/Vista 32/64-bit, 7), Symbian (S60 3rd Edition) and Linux-based operating systems. The universally-adaptable IPSec client includes FIPS Inside[21] certification and is seamlessly compatible with virtually any gateway on the market. The user-friendly GUI and intelligent policy enforcement provide even non-technical users with a one-click VPN. The client's latest features include Wireless Service Provider Roaming (WISPr) and streamlined Universal Mobile Telecommunications System (UMTS) Card support as well as a budget manager to manage wireless minutes and cost.

20 http://www.cisco.com/en/US/products/hw/vpndevc/ps2284/products_security_advisory09186a00805f0147.shtml

21 http://www.ncp-e.com/en/about-us/references.html

NCP engineering GmbH
Dombuehler Strasse 2
90449 Nuremberg
Phone: +49 911 99 68-0
Fax: +49 911 99 68-299

www.ncp-e.com

NCP engineering, Inc.
444 Castro Street, Suite 711
Mountain View, CA 94041
Phone: +1 (650) 316-6273
Fax: +1 (650) 251-4155

July 2010