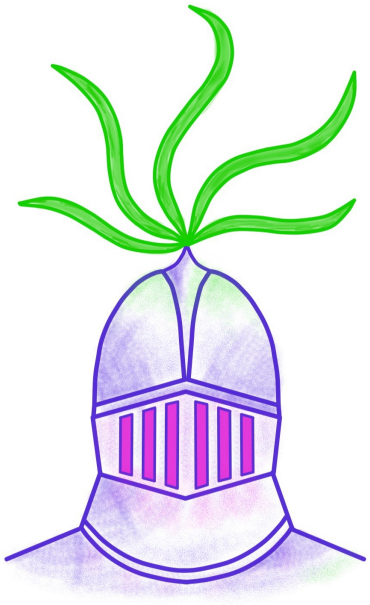

Network Security Design – New Considerations for 2015 and Beyond

September 2015 Whitepaper 2

Anthony Kirkham
tkirkham@neon-knight.net

www.neon-knight.net

Version: 1.01

Neon Knight Pty Ltd
Cybersecurity and Network Consulting

New Consideration for Network Security Design - 2015

Preamble

This is the second of a series of White Papers on Network Security Architecture.

I originally started writing this White Paper as I believed there was a need to examine the new and emerging requirements of a Network Security architecture. As I progressed, it quickly became clear that a White Paper covering the fundamentals was required ahead of this discussion. On that basis, I recommend reading the first White Paper prior to this second one.

This is a far from exhaustive coverage. It was written mostly while on planes, as a consultant, that pretty much a weekly occurrence.

I am focusing on a corporate world as I believe there is a great need to improve the security of these environments to enable business to operate in a secure low risk environment.

Introduction, Overview, Goals

Network security architectures are an area which have progressively evolved over the last 20 years. Over the same period, massive changes have occurred in both how the Internet and corporate networks are used, and the applications which operate on these networks. At the same time, massive changes have also occurred to the threat landscape facing these networks and the devices connected to them. While security technology has also made significant advances, it is my view that in many cases, the fundamental architectures deployed in many corporate networks have not evolved or been re-architected to the extent needed to deal with today's very different world.

The goal of this section is look at the security landscape today and the areas and trends which affect the design of Network Security Architectures.

Security Landscape Today

Firstly, it is necessary to have an understanding of the general classes of attacks today, their risk level and the approaches to mitigation. Figure 1 aims to provide a general representation of these considerations.

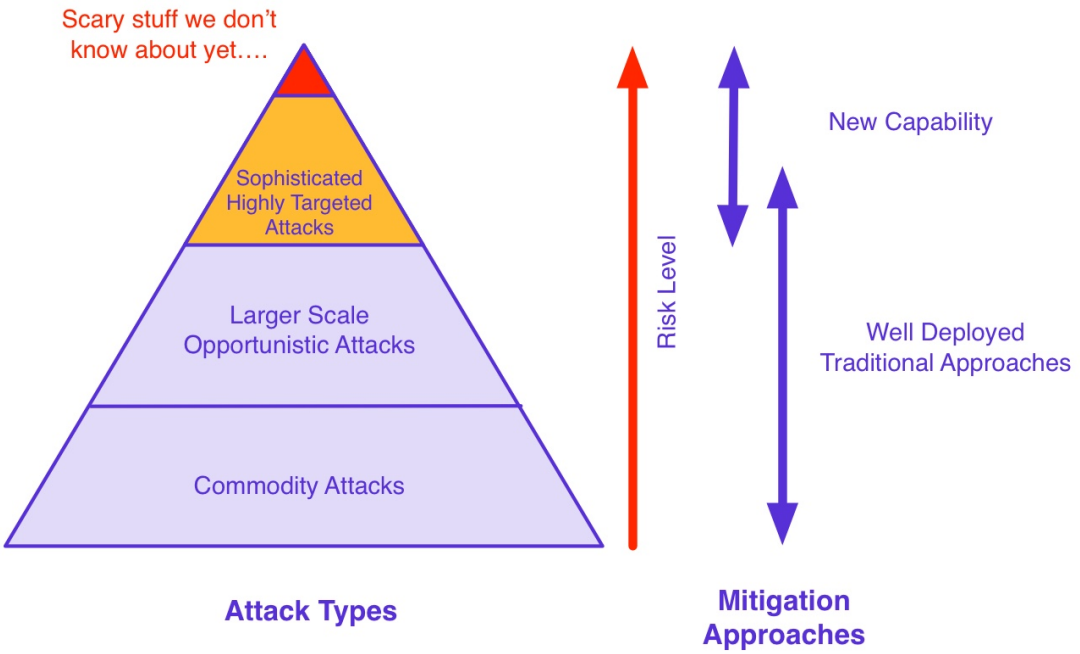

Figure 1 - Attack Types

While there is often a focus on Advanced Persistent Threats (APTs) and some of the sophisticated malware dangers, it needs to be remembered that a large percentage of the information security problem still comes from traditional and commodity attacks. There has been significant evidence presented in recent years of certain nation states utilising commodity attacks, but on a massive scale. In these cases, the attack sophistication has not been the main problem. It was clearly shown that these miscreants could work very effectively using only moderately sophisticated attack techniques at best. This means that well deployed traditional approaches are still important.

With that said, APTs and sophisticated malware is being used in targeted attacks today. With the use of Day-0 Malware in particular, traditional approaches will not protect from this class of attack, new capability is required.

What has not changed, or changed only a little?

First and foremost, a Network Security Architecture is a solution to address a Host and Server problem. It is a means of mitigating and controlling attacks on these end systems, controlling the spread of attacks if or when they are compromised. Additionally, it means having the operational capability to deal with compromises or sophisticated attacks, should they occur.

Host, Servers and End User, Mobile devices are still targets.

No surprises here. In general operating systems have been strengthened over the last decade. New versions of Windows, OSX and Linux have all added many fundamental changes which have made them harder to exploit and compromise.

However critical vulnerabilities are still being found in both the operating systems and the applications which run on them. New Zero day exploits can fetch many hundred of thousands of dollars, or a good number of BitCoins on the black market.

There are large populations of end user devices which are running outdated operating systems, or unpatched application or operating system versions. These devices are either highly vulnerable to exploitation or are already compromised. The sheer size of some Botnets is evidence of the magnitude of this situation and the number of compromised devices which exist in the wild. From a security design perspective, we must assume that some significant percentage of the user devices will be compromised at any point in time. These devices may well be connected to the inside of your corporate network and will likely try to initiate attacks within corporate networks. Network security design needs be performed on the premise.

While that may be the reality of end user devices, the role of host security is still unchanged. End user devices and servers should still be hardened in line with best practice. The same can be said for mobile devices. These should also still employ hardening, security configuration and mobile device management best practices.

Printers and Multifunction Devices are Particular Targets

Printers and Multifunction Devices are well known in the hacker communities for being unpatched and usually having well known and easily exploited vulnerabilities. They are often installed and forgotten about. Often these devices are leveraged within attacks and are regularly found host Command and Control Channels.

Network infrastructure is still a Target

Network infrastructure is still a high value target, except the methods of attacking it are more widely and better understood. Securing the network infrastructure is still critically important as compromised network infrastructure can be used to provide an excellent foundation to progress an attack once inside an organisation. A number of mature attack tools currently exist to perform network infrastructure attacks. For example, traffic traversing a router or switch can be redirected to a rogue system to intercept that traffic or snoop on the contents. With access to the network infrastructure attackers have the ability to bypass critical security controls, such as firewalls through manipulation of the network infrastructure. When an attacker has administrative access to the network infrastructure, it is potentially very easy to move into high security zones of the internal network.

In the domestic market, analysis of large Botnets has shown that very large numbers of commodity home router/firewall devices have been compromised are being used as the sources of Distributed DDoS attacks. There is a high chance

that in time these miscreants will realize they can use this compromised infrastructure to perform attacks like snooping on user traffic, or progressing attacks inside home networks.

The key point of this section is to emphasise the importance of a robustly secured network infrastructure. There are many good publications from the major networking vendors which describe in great detail, how to properly secure network (i.e. router/switch) infrastructure.

Application Design

Today application design often involves a number of servers interacting to form the overall application. For example, a dynamic web application will often use web server, some application logic and a database. As another example, web services may have many complex interactions occurring across many servers.

As a general observation, applications and application stacks are not written to interact with network security devices to produce the best security posture.

What has changed?

While there are many areas of the landscape which not changed very much, there are many areas which have. These following sections will discuss those in more detail. Later sections will provide some advice for dealing with these issues and considerations.

Complexity

Over the last decade just about everything has become more complicated. In particular, application architectures have become significantly more complex and more mission critical at the same time. More external business partners accessing systems are accessing systems from meshes or interconnected organisations.

This makes the job of securing these systems both harder and more complex.

Complexity is one of the worst enemies of security. Today, dealing with complexity needs to be viewed as subject area in itself. Highly complex systems present a number of risks including oversights in the design process, configuration errors, and deployment oversights.

Application Proliferation

There are way more applications on the network today. Many are legitimate business applications, many are personal apps, some are high risk and some are simply plain bad and have no place in a business environment.

When the Internet was first created, applications were based on an IANA assigned TCP/UDP port. Web used port 80, SMTP mail used port 25, DNS used 53. Today, many applications do not follow this process. For example, there are many hundreds of applications built on TCP 80, including social media, browser based file sharing, collaboration, and business systems. Additionally, many applications use dynamic port assignments or open secondary channels - Skype being one notable example. Many applications are also constructed based on a compound set of other protocols, Micrososft's Active Directory is one significant case in point. Many applications are highly networked with complex flows, such as peer-to-peer based applications. It is also very normal behavior for applications "call home" to check licensing and for available updates. Depending on the point-in-the-network, a different application mix can be expected. For example, the application mix that will traverse the Internet perimeter from the user population, will general be very different to the applications which are hosted in the data centre. In any case, trying to control application usage solely on port and protocol has become an unreliable and difficult approach.

Threat and Attack Landscape

Motives

Motivations have become way more dangerous - If we step back to the late 90' and early part of the 2000's, the motive for hackers, to a very significant degree, was notoriety. There was certainly an underground hacker community with this as their primary goal. We saw many of the early Internet Worms such as a Nimda, SQL Slammer and Blaster written to fulfil this goal.

From here, the motivation moved to that of profit. Theft of credit card numbers, and other personal information used for identity theft became significant targets. The miscreant community matured to a point where individuals provided specific specialist functions, each adding 'value' within the underground community, with a flow of revenue occurring between these parties.

In more recent times, the large scale theft of valuable intellectual property by Nation States has grown to be a massive problem. Observations have also shown that when organisations are compromised, the duration of the intrusions can be significant. Many of today's intrusions have been in place for many months and in some cases years, with information harvesting and leakage sometimes occurring on a continuous basis over these periods.

Today, additional and dangerous motives have begun to emerge from several nation states, where they have realised that an investment in Cyber warfare can provide a significant impact. These nation states have one goal - that is revenge and disruption to Western nations. Targets can include corporations, but of greater concern is specific attacks on Critical Infrastructure, Industrial Control Systems, many of which are highly vulnerable.

Security architecture needs to consider these motives and threats in the design process.

Client Side Attacks

One of the greatest changes in the security landscape in recent years has been the way in which organisations are compromised. Ten years ago attackers would scan a target organisation looking for open ports and vulnerable applications which could be exploited to gain entry. While this is still very much a valid attack technique, approximately 90% of successful compromises today are the result of client side attacks. The means users desktop PCs or laptops, with a vulnerable browser, are compromised by simply visiting an affected website. While 'questionable' or high risk web sites are a big issue here, there are many examples of where legitimate web sites have been compromised to host malware and attack clients visiting the site. Often these attacks are triggered through an initial Phishing email getting a user to visit a malicious or compromised site.

In addition, practices such as users running Java in the same browser that is used for Internet browsing is a high-risk activity. This practice can allow attackers via malicious web sites to compromise systems through vulnerable Java versions.

Day Zero Attacks - Signature based detection is becoming less effective

For many years Anti-Virus technology has been one of the foundation of protection against malware. It is also a mandatory element for organisations who require PCI compliance. While signature based AV is still an important element of any security strategy, it is becoming progressively less and less effective. The same goes for signature based IDS and IPS detection. Today the sheer number of malware variants and the rate at which they are being produced, coupled with Day 0 malware, is making it harder and harder for this approach to achieve the protection levels that were achieved in past.

With that said, there are still a huge number of very dangerous threats which can be stopped through the use of a quality IDS/IPS and Malware/AV scanning engine. But it is essential these systems are constantly maintained with current signature sets to be effective.

To be clear, signature based systems, be they IDS/IPS, AV engine, or network based malware detection/protection systems, will only detect known threats!

The detection of Day Zero or Unknown threats require the use of Sandboxing techniques where the questionable code is executed in a Sandbox and its behaviour analysed for malicious activity. A number of these systems now exist on the market from vendors such as Fireeye, Sourcefire (now part of Cisco Systems), Palo Alto Networks.

Malware delivery capability - Office Docs, PDFs, etc..

Another of the most significant changes which has occurred in the last few years has been the approach of miscreants to utilise new malware delivery mechanisms. Embedding Malware within office documents or PDFs has now become a somewhat widespread malware delivery technique. These techniques can take advantage of vulnerabilities in applications like the Office Suites, or PDF readers. Alternatively, office document macros, which many users leave enabled by default, can be used to execute the embedded malware.

In sophisticated or targeted attacks, Day-0 exploits or other malware embedded within one of these document formats is often used to gain the initial foothold into an organisation. Unless those documents are scanned and analysed at the perimeter, then the exploit or malware, will likely not be detected.

Widespread SSL/TLS usage

There are more and more applications using SSL and TLS encryption by default. Google, Facebook, just about every webmail application on the planet now uses SSL. Unfortunately, so does a lot of malware. Estimates vary from network to network, but a figure of around 30+% of web traffic utilising SSL/TLS is a good assumption to work on.

Once traffic is encapsulated within an encrypted SSL/TLS tunnel, its contents are virtually invisible to network security devices such as firewalls and IDS/IPS's. While some of the better IPSs can detect some malware tunneled in SSL from the packet size signatures and other heuristics, for the most part, once traffic is in SSL, you have no visibility of it. Additionally, SSL/TLS can be easily used as a Malware Command and Control (C&C), or a covert channel for information leakage out of an organisation. The same problems exist with other encrypted protocols such as IPSec, Secure Shell (SSH) or Secure Real-Time Transport (SRTP), however given the widespread use of SSL/TLS, it is currently by far the biggest consideration with security designs.

Virtualisation

Virtualisation is now prolific within data centre deployments. While this is an obvious statement, the impact it has on security architecture may not be quite as obvious.

The first consideration is that a virtualised environment requires the same architectural considerations and security controls as physical infrastructure. Many of these controls will need to be deployed within the virtual infrastructure, as opposed to relying on external physical devices. While security capabilities within virtualised environments have improved significantly in the last few years, the form they take is often different and from an industry perspective, the general understanding and operational experience is not as great as traditional approaches.

Some points to note on virtualised environment security:

From a network security design perspective, the same principals as the physical world still apply. As an example, I saw one organisation which had deployed nearly four thousand (4000) VMs within a single layer 2 network. Everything we have learned in the last 25 years of networking tells us that this is a really bad idea. What was really scary in this case was the vendor involved (or at least their local staff) endorsed the strategy and did not seem to have a grasp of the networking and security issues associated with such an approach.

Virtualisation can make aspects of Security Operations easier. For example the process of destroying a compromised VM and instantiating a new one is a far easier process than with a physical server.

Due to the agile nature of virtualised environments, remediation or re-engineering of security architectures has the potential to be easier. Designing and deploying a greenfield environment is always easier than reengineering of an existing environment. Virtualisation brings agility and and new suite of virtualised network security capabilities which can be leveraged within the environment.

There are now a number of mature Automation systems available for virtualised environments. Human errors and deployment inconsistencies have always been a significant security risk, for example failure to deploy a security control. The use of automation combined with prepackaged deployment patterns, including security profiles for VMs, can the reduce the potential for human errors.

From a risk perspective, the shared Hypervisor model is seen by the miscreant community as a high value target. Should a situation ever arise where the Hypervisor is successfully compromised, then all VMs running on that Hypervisor are potentially at risk. This risk should be understood for all virtualised environments. A general recommendation exists that only VMs of a like security affinity should be deployed on the same hypervisor. For example, deploying a public facing web server and a database server containing credit card information on a common Hypervisor, would probably be a bad idea.

IPv6

IPv6 is today available on just about every network device. If IPv6 network services are available, then IPv6 transport will be used in preference to IPv4. Many of the now current server operating systems will not allow IPv6 to be disabled without breaking things. From an architecture perspective, this means you can't just take the attitude "We don't use IPv6, so we don't need to worry about it".

IPv6 adds nothing to improve security. It only solves address space issues, with a few protocol improvements, nothing else.

While IPv6 deployment is still limited in most organisations, its security implications can not be ignored. Today, both Malware and miscreants utilises IPv6 in IPv4 tunneling techniques as covert channels. If you see either Protocol 41 or ISATAP traffic traversing your Internet perimeter, there is likely a very big problem. Minimally these type of issues need to be understood.

The generally accepted approach to IPv6 is the deployment of a dual stack environment, where both IPv4 and IPv6 run over a common network infrastructure. From a security perspective, it is critical that equivalent security controls are deployed for IPv6 as for IPv4. For example, if there is a traditional IPv4 firewall IPS or access list, they need to perform the same security function on the IPv6 transport as IPv4.

Mobile Devices

Just about everyone today has a personal mobile device, an iPhone, an iPad, an Android phone or tablet of some description. People are using their devices of their choice and applications of choice. Unfortunately, the whole mobile application space is awash with a huge number of Trojanised applications and malware. Application providers don't invest money in application development to simply give the apps away for free - there is always a catch, there is always a revenue stream somewhere in the deal.

How many people have corporate mail on their personal mobile device? Or access corporate applications or data from their personal tablet or iPad? How many staff are clueless when it comes to mobile device security practices? Certainly these mobile devices can bring significant productivity improvements, but from a security perspective it can means that your valuable corporate data can be accessed by a device population which may not controlled by any corporate IT policy. It also means that this data can possibly be access from places outside of the physical office environment.

Security architectures need to consider Mobile Devices and the remote access they provide as a distinct design consideration and not just as an add on. Mobile device access to needs to consider;

- The trust level of the mobile devices - does a corporate Mobile Device Management (MDM) capability exist?
- What happens if a device is lost, stolen, or left in a cab or a bar? Or left in an overseas hotel room unattended?
- The data they are allowed to access - do they require unrestricted access to potentially sensitive corporte information?
- The security technology used to access the corporate network. For example an IPsec or SSL based Remote Access VPN, or has some cruder less secure approach been cobbled together to provide access without a thorough understanding of the risks involved. For example, I have seen numerous deployments where a perimeter firewall port has simply been

opened from the Internet to the internal mail server to provide this access. Users are happy but its quick, dirty, dangerous and very common.

BYOD

There are any organisation who are adopting a Bring Your Own Device policy, also know by some as a Bring Your Own Disaster policy. In places like schools and universities, the use of personal devices on the network is very widespread. To a lesser extent business are adopting this policy, but the reality is that the prevalence of personal devices in the workplace is increasing. As such, the trust level of these user devices should be classified as 'low' given that these devices are not maintained under any sort of corporate IT policy.

If BYOD environment, consider;

- Utilising a 'Guest' network which provides only Internet access without any, or minimal access to corporate resources.
- One or more dedicated security zones for these devices, separate to the security zones where corporately managed devices reside.
- The use of Virtual Desktop Infrastructure (VDI) for access to corporate applications and resources.
- A Mobile Device Management (MDM) solution.

User Movement and Dynamic Addresses

Back in the early days of networking, users would get a DHCP address, typically for their desktop machine and would usually keep it for extended periods. These days users will move from one network location to another, often several times a day. For example from a wired desk connection, to a wireless connection in a meeting room to a VPN connection during a meeting at a coffee shop. Each of these network locations utilises a different IP address. The key point being that trying to identify users by their IP address is a very difficult process without the use of some good tools.

In addition, because of the dynamic nature of users IP Addresses, it is near impossible to apply any sort of network level security policy on a per user basis.

Back in the early days of networking, users would get a DHCP address, typically for their desktop machine and would usually keep it for extended periods. These days users will move from one network location to another, often several times a day. For example from a wired desk connection, to a wireless connection in a meeting room to a VPN connection during a meeting at a coffee shop. Each of these network locations utilises a different IP address. The key point being that trying to identify users by their IP address is a very difficult process without the use of some good tools.

In addition, because of the dynamic nature of users IP Addresses, it is near impossible to apply any sort of network level security policy on a per user basis.

The Cloud

Cloud Computing and services have become a highly economical and attractive option for many organisations. From a security architecture perspective, the situation is much the same as for virtualised environments - the same security constructs need to be deployed. Most of the cloud providers do a very good job of offering commoditised on-demand compute and storage capability (i.e. scaling wide) but in most cases the internal security constructs like virtual LANs, virtual routers and other virtualised security controls are not available or are in the early stages of maturity. From an implementation perspective, to build anything beyond a simple enclave (such as a secure n-tier architecture) can be difficult.

At a technical level, Public Cloud utilises a shared Hypervisor model with the same potential risks as described for virtualised environments. In some case, this won't present a significant business risk. However, in other cases the risk may be completely unacceptable. This is something that needs to be evaluated on a case-by-case basis.

On the flip side, many of the cloud services are implemented on well secured network and virtualisation infrastructure. This means that organisations can avoid both, the need to build the infrastructure and also the need to properly secure that infrastructure. The later is often a non-trivial task on which many organisation have not done a good job. While many security professionals often view Cloud as fought with security risks, which in many cases is entirely justified, there are some significant benefits for some use cases. With all that said, offerings do vary widely and the use of cloud services and the vendors should be carefully assessed.

Technology Advances

In the preceding White Paper, I mentioned that constantly monitoring new technologies is an important activity. Over the years, many vendors who have had dominant market positions have lost their market share to newer and more agile companies. While I have discussed the changes to the Threat Landscape, there have also been significant technological advances to address many of these new threats.

Even in first half of 2015, despite some 'major hack' making the news on a daily basis, I hear people arguing to keep the status quo. I have heard intelligent people argue to "keep what we have .. as that is how we have always done it". I'm not advocating moving to new technology for technology sake as it is essential to understand both the problem and how a new technology can address it. But I certainly don't believe it is a healthy situation to allow this situation.

In particular, there have been significant advances in security technologies, for example;

- Security Analytics
- Next Generation Firewall and IPS technology
- Firewall security policy management and auditing tools
- Anti-exploitation technology - I am closely monitoring both Cylance Protect, SentinalOne, and Traps from Palo Alto Networks.

Social changes

Social media is huge and in wide spread use within the workplace with the number of social media sites continuing to increase. While a lot of Social media is personal use, today it is also an essential business and marketing tool. So it's not just as simple as closing it off. From a business risk perspective, it becomes very easy for staff to inadvertently leak confidential information onto social networks. Some social networks are also vectors for social engineering attacks and conduits for the distribution of malware.

Today, a corporate Information Security Policy needs to consider acceptable social media usage in the workplace. While this is an issue which will vary widely from organisation to organisation, an example may be:

- General staff have unrestricted access to read and post to Facebook, LinkedIn and Twitter.
- The use of Facebook or other social media applications and games is forbidden.
- The use Tumblr and Instagram is forbidden.
- Marketing staff can access Facebook and LinkedIn for the purpose of updating social media based marketing information.
- Corporate communications staff can access Twitter for the purpose of providing updates via this media.

Every organisation needs to have at least an initial policy and have performed some risk assessment associated with social media usage in their environment. An Internet Perimeter Architecture today needs the capability to monitor and control social media applications to support that corporate policy. There are a number of technologies which can perform this function including Next-Generation Firewalls, IPSs, and Web Gateways.

Hardware attacks and Compromise

Traditionally when we think of Malware, we think of a malicious piece of software, a Trojan, or an exploit which compromises another system or an operating system. While this is the case in the vast majority of situations, looking forward we need to also consider hardware compromises. Today, many components within commodity hardware systems contain a reasonable level of intelligence. For example, many peripheral interface chipsets on PCs, MACs, and Server Hardware contain an embedded Linux kernel. BIOS chipsets usually contain flash memory to enable an easy upgrade process.

Given that this hardware is accessed through software based device drivers, it presents a significant attack surface. This potentially allows Malware to compromise these peripheral devices and stay persistent even after a fresh installation of the operating system and application set. On this basis, we must assume that these type of hardware compromises won't be detected through the use of traditional AV and should be considered as Unknown Malware.

Additionally, there is also a risk that malicious capability be loaded into hardware during the manufacture or supply chain process. I'm not going to name them specifically, but there have been recent cases of this situation occurring. The problem is potentially very real.

Architecturally, these risk add to the case for Analytics and network level behavioral analysis. For example, if a system or systems are sending questionable traffic to a foreign destination, then you have a situation which requires investigation and explanation.

The Importance of Operational Capability

Today, continuous monitoring and security operations capability are critical. Many reports and analysis of breaches has shown that attackers who have successfully compromised organisations, have had a presence within the network for extended periods without detection. A number of industry reports have also clearly shown that a high percentage of breaches are not detected by the compromised organisation itself, but through third parties who have detected the compromise through external means.

Relying on a yearly penetration test to detect a breach is a completely insufficient approach. Either In-House and/or outsourced security monitoring is a necessity. The underlying security architecture must be designed to easily enable this activity.

Architectural Recommendations –2015 and Beyond

Introduction

In the first half of this White Paper, I have discussed a number of macro trends and changes to the security landscape which have implication for Security Architecture and Network Security Architecture. In some of these discussion topics, I have made brief recommendations. The second half will build on those further by providing more in depth discussion around a number of design considerations and recommendations which I now view as essential capabilities or elements of an effective Network Security Architecture.

I will break this into two parts;

- Architectural Foundations
- Operational Enablement

Prior to these two sections, I wish to introduce the Lockheed Martin Cyber Kill Chain[1] which will serve as a model to help frame both the Architectural Foundations and the Operational Enablement discussions.

Introducing the Cyber Kill Chain

A recent approach being adopted by a number of thought leading organisations is the Cyber Kill Chain., which was developed by the Lockheed Martin Corporation. The approach is documented at:

http://www.lockheedmartin.com.au/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf

The first element of the Cyber Kill Chain is a model which deconstructs and describes an attack lifecycle. Within that context it can be used for framing an analysis and incident response capability, particularly where advanced or targeted threats are involved, such as APTs.

The purpose of this White Paper is to discuss the new requirements of security architectures. Using the Cyber Kill Chain model, we can ensure where possible, the architecture provides the necessary underlying capabilities to detect and respond to an advanced attack at each stage of the lifecycle. Figure 2illustrates the seven stages of the attack lifecycle as described in the Lockheed Martin Cyber Kill Chain model.

[1] CYBER KILL CHAIN is a registered trademark of Lockheed Martin Corporation.

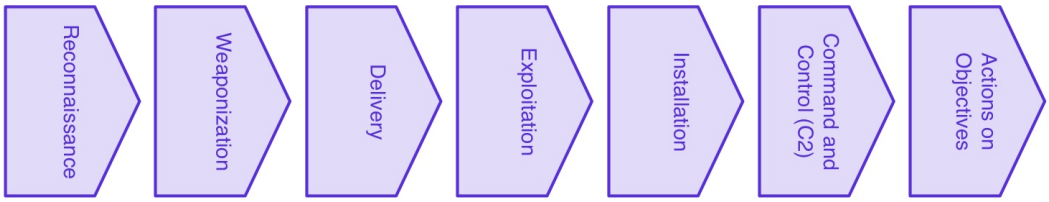

Figure 2 – Attack Lifecycle - Seven Stages

Table 1 provides a description of each of those seven stages. I have added Step 6a to describe the process of lateral movement. I believe it is important to distinctly define this stage as it has particular relevance from an architectural perspective.

1	Reconnaissance	The process by which an attacker researches, identifies, selects targets and attack vectors.
2	Weaponization	The creation of an attack tool (i.e. weapon), possibly customised, usually by coupling a remote access trojan with an exploit into a deliverable payload. Increasingly, client application data files such as Adobe Portable Document Format (PDF) or Microsoft Office documents serve as the weaponized deliverable. Increasingly 'Day Zero' exploits and malware are being employed, in particular for 'high value' targets.
3	Delivery	The transmission of the weapon to the targeted environment. Typical mechanisms include Email attachments (i.e. Spear Phishing), malicious websites, and removable USB media.
4	Exploitation	After the weapon is successfully delivered to a victim host, an exploitation process is used to trigger the intruders code. The exploitation process often targets an application or operating system vulnerability, but can also simply exploit the users themselves or leverage an operating system feature that auto-executes code (i.e. Microsoft Word Macros).
5	Installation	Installation of a remote access trojan or backdoor on the victim system allows the attacker to establish persistence inside the environment.
6	Command and Control (C2)	Typically, compromised hosts will beacon outbound to an Internet control server to establish a C2 channel. Once a C2 channel is established, the attackers have an internal presence within the target environment.
6a[2]	Lateral Movement, Internal Attack Propagation,	Once an initial foothold is established within the environment, the attack must progress internally. Often steps 1-5, or variations of them are repeated

[2] Step 6a is not a part of the original Lockheed Martin Cyber Kill Chain.

	Search for target data.	for this process. A key goal is to search for, or locate, either 'interesting' or 'target' data.
7	Actions on Objectives	Typically the objective will be data exfiltration which involves collecting, encrypting and extracting information from the victim environment through some transmission mechanism. Alternatively, violations of data integrity or availability can also be potential objectives.

Table 1 - Deconstructed Attack Lifecycle

An understanding of the above process allows us to design appropriate security controls at each stage of the attack life cycle. While the Cyber Kill Chain does an excellent job of articulating that process, it additionally describes others actions, concepts and capabilities which are necessary to quickly identify an advanced attack and engage an appropriate response.

In this day and age attacks are not just single events, but campaigns, which will typically use multiple attacks on multiple vectors or targets in the attempt to achieve the objectives. Both defensive security design and operational security models need to think on this basis.

The Cyber Kill Chain is based on the paradigm that no network is free from compromise, despite all best defensive architectural practices. To provide operational enablement under this paradigm requires detailed monitoring for Indicators of Compromise (IOC), Indicators of Attack (IOA) and any other anomalous behaviors. A Security Architecture, coupled with appropriate analytical tools must now provide these capabilities to allow effective operational enablement.

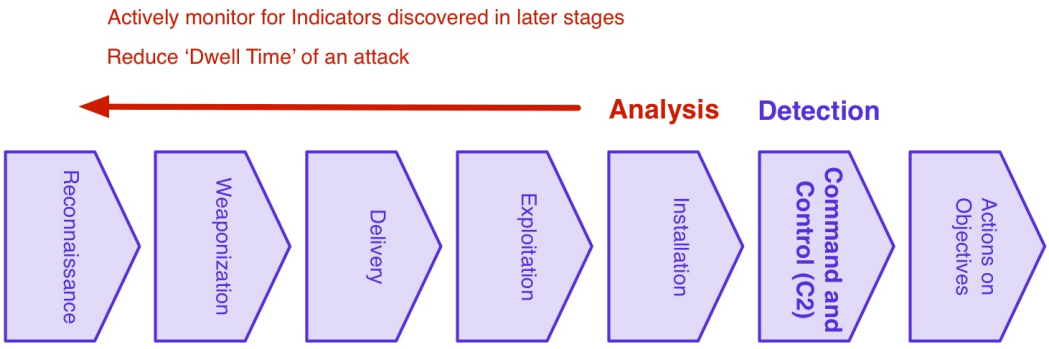

Figure 3 - Enabling Early Stage Detection

From an operational perspective, this paradigm can utilise intelligence and indicators gathered in later stages of an attack campaign (i.e. IOC, IOA) to enable earlier detection and response. This concept is illustrated in Figure 3. This approach can also be known as "Shortening the Kill Chain" and allows operations staff to detect and block threats at earlier stages with speed and confidence. Or, in other words, reducing the Dwell Time, where an undiscovered attack is operational in the network. From a cost perspective, it is significantly (to

massively) cheaper to block an attack at the earlier stages, than to deal with the consequences of a later stage breach.

Living Off the Land Attacks

Within the deconstructed attack lifecycle I specifically called out the Lateral Movement stage. Often this stage provides some of the most significant insight into attack detection.

While Malware and a variety of attack tools are still very commonly used to move laterally within an organisation, there has been a new trend noted where attackers are making use of the generally available IT tools and systems for this purpose. Given that these tools are legitimate software, the misuse of them becomes much harder to detect.

This trend highlights the need for sophisticated tools and analytics to detect these misuse scenarios.

Architectural Foundations

Beyond Defence in Depth

Going back 10 or 15 years, the M&M analogy was often used to describe many networks - "crunchy on the outside, soft on the inside", meaning there was usually a hard perimeter with very little in terms of security controls on the internal network. There are still many networks today which are still based on this very outdated approach, or have never moved beyond it.

Figure 4 - M&M Analogy

As was described above in the CKC model, successful cyber intrusions involve a number of steps. While it is one thing for a attacker to gain a foothold internally, a number of other steps are normally required for the attacker to successfully complete the intrusion and achieve their objectives. Usually this means exfiltrating information from the organisation without being detected.

When several recent, successful, high profile, Cyber Attacks were analysed, it was noted that in most cases the highly sophisticated element of the attack was only used in one stage. Then generally more commoditised tools were used in subsequent stages (such as lateral movement), or to complete the attack.

In this day and age, an M&M architecture is completely insufficient. From an architectural perspective ensuring that appropriate security controls are used to prevent, or at least have visibility of attacks at each stage of the CKC is an approach which should be adopted. In saying this, naturally the cost versus the benefit needs to be a design decision. However, if this approach is utilised, it forces the attackers to require sophisticated techniques (such as a Day Zero, or near Day Zero) at every stage for it to be successful. While this is not an absolute guarantee of protection, it forces the attackers to expend more time and effort, hopefully making it uneconomic for them to invest time in your organisation as a target.

For a Security Architecture to now be effective, it should aim to use strategically located security controls and techniques so as to block intrusions at each step of the process. If we work on the assumption that we are dealing with very

sophisticated attacks, security architectures should work on the principal that if the first stage of the attack is successful, then internal defences are constructed to prevent subsequent stages of the attack from being successful. For example, requiring the attacker to utilise a different, sophisticated attack technique (such as a Day Zero, or near Day Zero) at each stage.

So if we work on the principal that in the event that the first stage of the attack is successful, other internal defences can prevent subsequent stages of the attack from being successful.

I think of this approach as taking Defence in Depth to a new level!

Managing Complexity

The management of complexity is a topic which could fill an entire text book. Given that the focus here is security, I will limit the scope to some general approaches which should be used within network security design. But first and foremost security engineers should understand that security solutions, especially large scale ones, are often highly complex deployments and need to be approached as such.

Major Issue One - it is essential that organisations understand what they have. It is amazing how many organisations have network and security infrastructure which is inadequately, poorly or simply not documented. This also extends to the management, structure and documentation of firewall rules and other security controls. Security is as much about process as anything. This is one area where a strict and mature process needs to exist and be followed. If you don't accurately understand your systems and their deployment, then it is impossible to manage it.

One of the most significant design principals for complex infrastructure, is the use of Modular Design. Modular Design is the process of breaking the problem into smaller repeatable blocks. Sometime these are known as design patterns. As an example, say if you are managing a large environment with many smaller branch offices. It would be prudent to utilise a single design and deployment pattern which included a common set of security controls, used at all sites. In the event that a security control, such as a firewall rule or access list required updating, then this change can then be easily applied to all sites.

As the complexity trend continues, I believe a modular design approach will become vital. Any large scale enterprise network based on bespoke configurations will either be unmanageable or have a very high operational expenditure associated with it. In particular, Firewall rules are often one of the highest maintenance areas. The use of modular blocks (as far as possible) within a rules deployment can improve the manageability, lower the operational cost and significantly improve the overall security posture.

For large complex deployments adequate tool sets are essential. Specific to security there are many tasks which now require tools to be managed accurately and efficiently. One of the prime examples is the management of firewall security policy. I have personally performed many security assessments for many organisation across the globe and I am yet to find one which did not have major problems in the management of firewall security policy. In this day and age, for most organisations, this task is now sufficiently complex that a quality tool is required. Management through manual processes is both error prone and time consuming.

When you consider a large environment of today, there are usually many sites, all interconnected, often many DCs, each with sometimes thousands of servers, often with many open ports on each, some of those systems will have exploitable vulnerabilities. Then you have access list and firewall rules to protect it all. Quickly you can see this is N to the power-of-something complexity problem. Tool sets are now essential!

Security Policy as far as possible should be organised in a Hierarchical Structure. Overwriting objects with site specific information can be used to keep a consistent structure.

Abstraction is the technique for hiding the lower level complexity of an implementation from the view above. Abstraction is an essential technique within software engineering. Code is written in blocks with Function Calls or Methods used to hide to underlying complexity. When designing complex security solutions, the same mindset should be used.

Zoned Architectures

In the first White Paper I spent a fair amount of time discussing this topic as it is a foundation architectural concept. I would recommend that you have read that material ahead of this discussion.

Zoned security architecture is as important today as ever. Historically, the concept dates back to perimeter firewall architectures which contained an Inside Network, an Outside Network (or the Internet), and a DMZ (Demilitarised zone). In more recent times, John Kindervag of Forrester Research released the concept of a "Zero Trust Architecture". Personally I dislike the term "Zero Trust" as I believe it is a quite misleading marketing term - there will always be some level of trust involved and if there wasn't, the implementation cost would be prohibitive. However, the paper does go onto the concept of "never trust, always verify" which is believe is a far more realistic approach. While I don't believe the paper introduced anything conceptually new, it does highlight the importance of the zoned architecture concept. It does however strongly emphasise the increased capabilities that need to be implemented for traffic traversing security zone boundaries. An architectural model to support the Cyber Kill Chain aligns with this approach.

I believe will be critically important in building effective security architectures for the present and near future.

What this means. Traditionally firewalls have provided a Policy Enforcement function, usually based on the Five-Tuple concept (Source/Dest IP, Protocol, Source/Dest port). As attacks are growing in sophistication, going forward the capability needs to include deeper inspection to limit the propagation of attacks and visibility which can feed into analytics systems.

Client Side Attacks - I noted earlier that there has been a massive shift in attack trends to where internal user devices are attacked and compromised through Client Side Attacks. It must be assumed there will be some level of compromise of user devices and as such a layer of protection of internal application infrastructure is critical. In addition to a robust perimeter architecture, an appropriately compartmentalised internal architecture is more important than ever. In practice this means at least Firewall, IPS and Day-Zero malware analysis in front of key business systems or applications.

Visualising and Managing Applications and Application Traffic

Previously I had highlighted the proliferation of applications. Some are essential business applications, others have no place in a business environment. Application visibility and control, in my opinion, are now an essential element of a security architecture. This is now mature capability which is available from a number of Next Generation Firewall and IPS vendors.

The management of application level traffic is highly dependent on the location in the network, also known as a Point-In-the-Network (PIN). Figure 5 illustrates this concept. Traffic which may be perfectly acceptable crossing an Internet perimeter, may be totally inappropriate traversing the Data Centre perimeter and vice versa. Typically very different application sets will be seen at these two point-in-the-network. Take VNC or another Remote Desktop Protocol, seeing this type of traffic going to a server in the data centre would probably represent a valid use case, such as server administration. However, if we saw the same traffic traversing the Internet perimeter, then it would likely be cause for alarm. Very different application sets will be seen at different Points-In-the-Network. This is critical to understand.

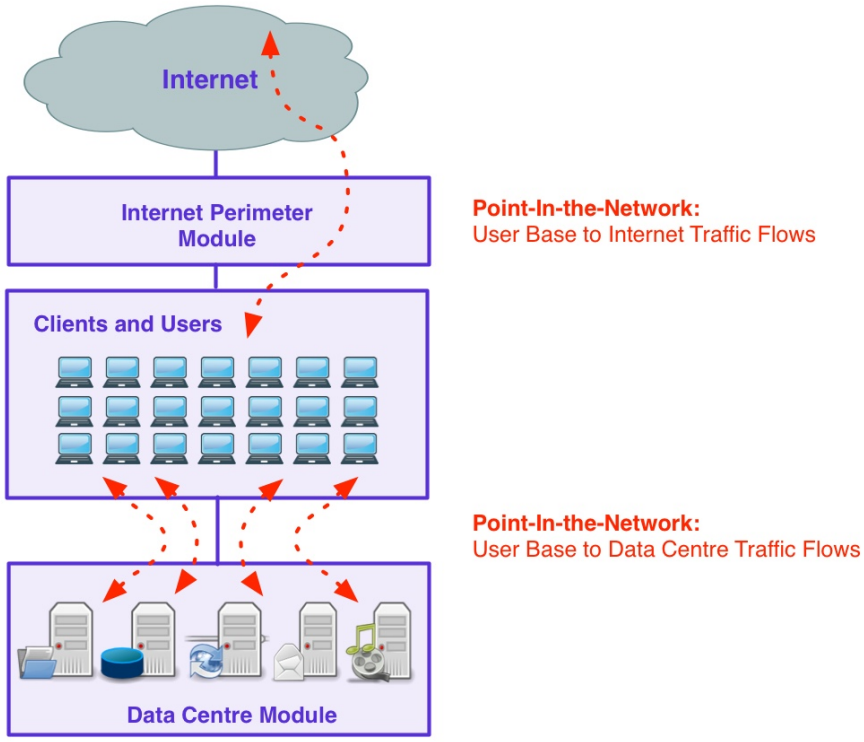

Figure 5 - Application Flows - Points-In-the-Network

As a first step, having application level visibility is a key capability. When you start seeing applications like TOR crossing your corporate Internet Perimeter, or BitTorrent consuming 40% of the expensive corporate Internet bandwidth for downloading pirated movies, you can quickly see the value.

It is from a Data Centres that corporate applications are hosted. Minimising the attack surface to these applications will significantly improve the overall security posture. As I say above the application flows at this PIN will be very different from the Internet Perimeter and also very different from and Business-to-Business flows which also need to be very tightly controlled and monitored.

When you begin to understand the application profiles, you can then move to controlling it through application level security policy.

While application control is a very worthwhile goal, which can dramatically improve a security posture, note that it can be a complex undertaking. Often applications are built top of one and other and non-obvious application dependencies do exist.

Netflow Based Visibility

Previously we have discussed the fact Firewalls and IDS/IPS devices can provide a very rich set of logging information at security zone boundaries. While the boundaries between security zones provide a very logical Point-in-the-Network for visibility, they are not necessarily the only places.

In White Paper One I discussed the use of Netflow (and S-Flow) for flow visibility within a network infrastructure. Netflow is a capability which is a fairly standard feature of most network level routers, switches as well as some Firewalls. I see the use of full-flow Netflow as a key foundational element of analytics solutions going forward. Capturing and storing full-flow Netflow records provides a full history of all conversations which have occurred through strategic Points-In-the-Network. This is analogous to a phone company keeping records of all phone calls (but without recording the contents of each call of course).

Netflow can provide an excellent means of understanding the traffic flows which are occurring both within your network and externally. By understanding which systems are communicating, and in what volumes, this information can be used as an input into network and network security architecture, particularly when reengineering a zoned security architecture. Coupled with Threat and Security Intelligence feeds, Netflow information can also be used to quickly identify internal devices speaking to questionable or malicious external sources.

Such a solution is illustrated in Figure 6. In this figure I have illustrated the Internet Perimeter and the Data Centre perimeter as the monitoring points. While are strategically important points, there are certainly others.

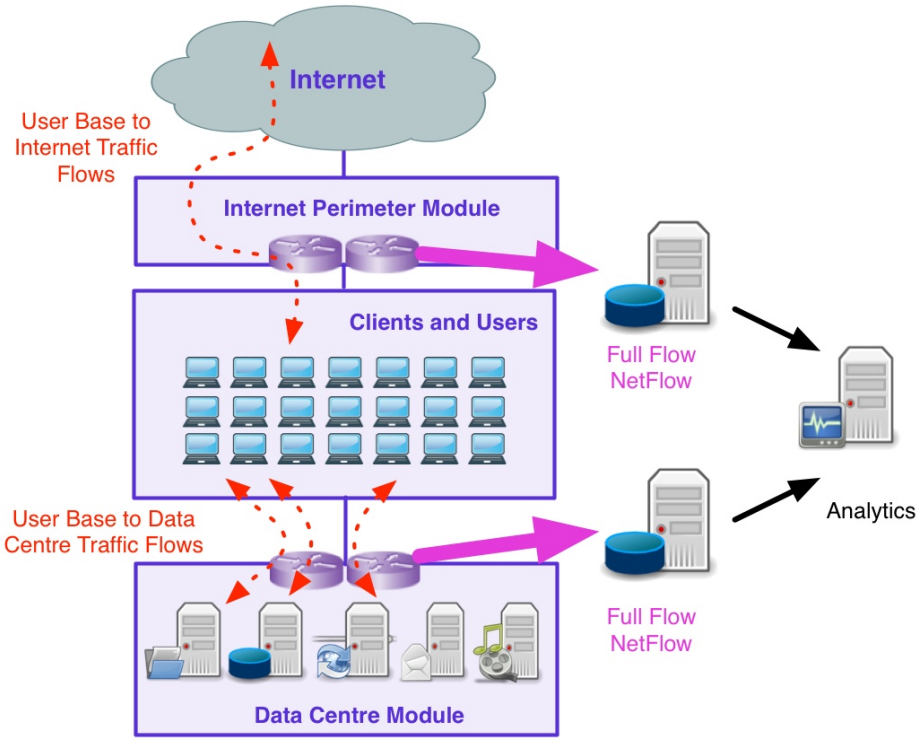

Figure 6 - Netflow based Analytics

With this approach presented, note that the storage of Netflow information will require a modest investment. Network speed, storage duration, and the number of monitoring points are all factors in the storage requirements which must be considered on a case-by-case basis. With that said, Netflow can provide a very significant capability within the CKC model. For example, should a new Indicator

be uncovered, a search can be conducted through historical Netflow records to identify if the other devices or network elements have been affected by the same attack campaign. Locate them, go back and remediate.

Packet Capture

While Netflow provides an excellent means of determining traffic flows within a network, packet capture can provide a far richer set of information, but with the additional cost of significantly increased storage requirements. Several Analytics solutions today can consume packet capture from strategic Points-In-the-Network.

In addition to enabling Analytics, continuous capture also enables a rich forensics capability. For example, I know of organisations who run a continuous stream packet capture of all traffic traversing the Internet Perimeter. This traffic is stored in a circular manner for a limited period. In the event of an incident, a full copy of all traffic is available to security operations or for forensics purposes.

Given the storage requirements associated with Packet Capture, the monitoring points need to be chosen very carefully. Typically this would mean in front of key information assets or servers, but downstream of any network security controls such as firewalls, IPS, Access Lists, etc. This eliminates capturing any crud and focuses packet capture on traffic which actually makes it to the server/servers, etc.

Ad-Hock Packet Capture provides a somewhat different capability. Ad-Hock Packet Capture provides the ability to perform an on-demand capture from a wide range of Points-in-the-Network. This capability provides both a valuable troubleshooting and security operations tool. Today most network infrastructures have the ability to mirror traffic from various PINs through the use of an Encapsulated Remote monitoring approach (i.e. such as Cisco's ERSPAN). Traffic stream information can be captured with a variety of tools (i.e. Pcap, etc) Today, it is particularly important this capability is available within any virtualised infrastructure.

Ad-Hock Packet Capture can be enabled for a very small cost and as such is a capability which should now be included in every Network or Network Security Architecture.

Visibility and Policy Enforcement within SSL and TLS

As I discussed earlier, SSL and TLS are in widespread use. Once traffic is encrypted within these tunnels, visibility of the traffic and possible threats, is gone.

Before I go too much further it is worth pointing out that SSL is now considered insecure and as a result is a deprecated protocol which should not be used. It has been replaced by TLS, currently version 1.2, which should be used in its place.

The pervasiveness of these protocols means having a capability to decrypt SSL/TLS to gain visibility into these encrypted streams, is a really big design consideration.

At the Internet Perimeter, outside of SSL/TLS, and possibly SSH (Secure Shell) for very specific purposes, there should be no other encrypted tunnels in use! Any other encrypted tunnel, without a very compelling and specific use case should be blocked. Virtually all Next-Generation Firewall Technology, or Network Proxies can do this. Additionally the use of TLS over the now deprecated SSL for all outbound web traffic should also be enforced at the Internet Perimeter. This design approach utilises TLS as the sanctioned encrypted tunnel, and allows decryption technology to gain visibility into the contents of the encrypted tunnel. SSL/TLS on non-standard ports would also be blocked without a specific use case.

A little background on SSL/TLS decryption, as two variants exist to service two use cases;

> Forward Proxies - This variant provides a means of decrypting SSL/TLS traffic initiated by a client, typically on the Internal corporate network, to an external server, typically out on the Internet. In this use case you have no access to the Digital Certificates in use on the destination servers. As no access to the Digital Certificates, and hence the encryption keys of the destination sever, a man-in-middle process is used to intercept the SSL/TLS traffic stream and decrypt it. This requires a subordinate CA on the interception device, typically a Next-Generation FW or a Proxy, with a corresponding certificate deployed on every client device. In a large network with many devices, this can be a significant task. Approaches like Active Directory Group Policy can be used to deploy these certificates to user devices, but the key point is that a degree of deployment complexity exists, particularly to non-Microsoft devices.
>
> Inbound Decryption - This solution caters for the inspection of SSL/TLS traffic destined for a Web Server for which you control and hence have access to the Digital Certificates deployed on those servers. This means servers within the Data Centre and any Internet facing Web Servers within the corporate DMZ (assuming these are under your control). This approach works by placing a copy the web server certificate on the SSL/TLS interception device. This is typically a straight forward process and allows visibility into SSL/TLS traffic streams.

Given the pervasiveness of SSL and TLS traffic today, visibility inside these streams is an essential architectural capability. It should rank among the most important design considerations. Next Generation Firewall and Next Generation IPS devices will typically provide visibility of applications, exploits and malware

(virus) files within the stream. Network Proxies are typically limited to virus file with some having IPS module capability.

From a product selection perspective, it is important to get accurate performance information about the solution elements, including both stream decryption performance as well as the transaction performance of the initial key exchanges. Some product utilise embedded hardware based crypto processor which provide excellent performance, other are more limited through purely software based deployments.

Managing New Attack Trends - Day Zero and Embedded Malware

Earlier I outlined the issue of Day Zero vulnerabilities and exploits, coupled with the newer Malware delivery trends. The problem is very real and having a capability to inspect files in transit for potential malware is now a necessity. This means performing both File Hash and Dynamic Analysis on problematic file types including executables, DLLs, office docs, PDFs, Java runtimes, etc.

> File Hash analysis simply takes a hash of the file while in transit and through a dynamic lookup will check to see if the hash has been previously identified as malware. A number of vendors are now using global Malware databases such that when a new piece of Malware is identified, the file disposition information is instantly available on a global basis.
>
> Dynamic Analysis is the process of executing, or in the case of office documents, PDFs, etc, opening the file, to analyse its behavior. Dynamic Analysis works on the basis that legitimate software behaves very differently to Malware. Behavioral analysis can usually quickly and accurately make a determination.

Architecturally, file inspection must be performed at the Internet Perimeter, with it also being highly desirable at the Data Centre perimeter. However the later will likely have significant performance considerations which will need to be assessed on a case-by-case basis.

All file transmission vectors should be analysed. This will include email, HTTP/HTTPS file downloads, FTP/SFTP, SSH/SCP, etc. Any vectors outside of these which can't be adequately inspected, should be blocked. In particular any questionable application protocols such as Peer-to-Peer application, should no be permitted unless a valid use case exists.

To provide an effective operational capability that aligns with Cyber Kill Chain model, Day Zero and Near Day Zero detection and analysis is a crucial element. In fact, the previous referenced Lockheed Martin Paper cites examples of where Day Zero malware had been used in attack campaigns against them.

Before I go any further, I will make one very important point. While Day Zero detection and analysis capability is important, don't forget about the fundamental security architecture principals! I can't emphasise this strongly enough. Many sources have clearly illustrated that the vast majority attacks are still utilising commodity tools and techniques. It is very easy to get caught up in the hype and fear around Day Zero threats and loose sight of a far larger risk area.

Security within Virtualised Environments

Security within virtualised environments is a sufficiently large topic to fill its own White Paper. There are really two key points:

The same security controls applied within a physical environment, should be applied within a virtualised environment. These days a number of vendors can provide virtual form factor Firewalls, IPS, Next-Generation Firewalls, switches and routers, etc. These can all be used in the same manner as their physical counterparts.

Hypervisor exploitation is an area of great interest to the hacker community. Going forward we should work on the assumption that at some point, a hypervisor vulnerability be found potentially giving an attacker a bridge between systems running on the a common Hypervisor. On this basis it is prudent to only keep systems of a like security affinity on the same hypervisor. Or in other words, don't place systems with significantly varying security requirements on a common hypervisor, or if you do, clearly understand the risk.

DNS

DNS Security is a subject area in itself, something I intend to write on in detail at a subsequent point. I'll briefly make two points:

Separate DNS infrastructure should be deployed for;

- Internet Accessible Authoritative Name Resolution. i.e. People on the Internet resolving your own namespace and servers.
- Internal Recursive Name Resolution. i.e. Your internal users performing name resolution of systems on the Internet.

Monitoring of DNS queries is a very valuable source of instrumentation which can provide standalone intelligence, which can subsequently be fed into analytics systems. Compromised systems will often try to access questionable domains on the Internet. This can be very quickly identified through their DNS queries.

Active Directory

Active Directory is an area that makes me really nervous. If you are a hacker, who has just breached the perimeter of any reasonable size organisation, Active

Directory and in particular the Domain Controllers, are highly likely to be your next target.

Active Directory will normally provide a significant number of services within most large organisations. Active Directory not only Authenticates and Authorises all users and computers within a Domain, but also contains a wealth of information on the various services within the Domain and where they are located. So, if your the hacker who has just managed to successfully breach the network perimeter, and your looking for, say, the database that contains the credit card information, then access to AD to go searching for things like SQL Services, is a highly valuable asset.

As I noted earlier, many Next Generation Firewalls are making use of AD to provide User Information on which firewall security policy can be based. This is another key function which increases the criticality of a highly secure AD deployment.

Active Directory infrastructure is often not controlled by the Security team. I see in many (or most) organisations Domain Controllers, being directly accessible to the user population. Often there are quite valid reasons for this. Active Directory is a compound protocol which utilises a wide range of both TCP and UDP ports. It is a difficult protocol to put through a firewall as many ports need to be open. In any case, the large number of open ports and services in AD provides a wide attack surface should vulnerabilities exist.

Given the criticality of AD, I believe the overall security surrounding this system needs to be given far more attention. While I am focusing on Network Security in this White Paper, a whole range of AD security controls need to be included to meet this goal. These include;

- Maintaining Operating System patch levels
- Maintaining Application patch levels
- Utilising current OS and software versions
- Making use of the newer security features of recent versions
- Application White listing on DC
- Prudent use of privileged accounts such as Domain Admin
- Security monitoring of high privilege activities

With that said, we can move the discussion back to the network security techniques which should be used to protect AD Infrastructure. While I stated above that AD requires a large number of ports to be open, this requirement does not mean that a significant security posture gain can not be achieved from the placement of Domain Controllers (DC) within a firewall protected enclave.

- Utilise either a Jump Host or a dedicated, single purpose, administrative workstation for privileged activities. Utilise FW based security policy to enforce this level of access.

- Restrict or block access from the DCs to the Internet. Apart from perhaps accessing software or patch updates (and even then this is questionable), there is no reason DCs should be accessing the Internet.
- Minimise the attack surface of DCs from user devices.

Not only can a firewall limit inbound access the DC, they can also block outbound traffic from the DC.

Appropriate FW protection must coupled with a quality IPS to prevent any exploitation of vulnerabilities should they exist or become known. In the event a new vulnerability is announced or disclosed, it is a lot easier to use an IPS as a compensating control, until a convenient time to patch the vulnerability is available. While FWs and IPSs predominately provide a Policy Enforcement function, they also provide very granular logging which enables visibility of the transactions. Within the next generation of security architecture, underlying visibility to enable analytics is a key capability.

It should also be noted that some Next Generation Firewalls (such as Palo Alto Networks) provide the ability to specifically identify and filter on AD as an application. While this isn't necessarily as easy to setup as the brochure suggests, it is a very valuable capability for protecting AD environments.

My key point here is that AD Environments are key high value targets which require a high degree of attention. Lessons from the recent Target breach clearly suggested that if diligent monitoring of the AD environment had been performed, then the abnormal activity would almost certainly have been noticed.

Blocking leakage of key information assets

The seventh stage of the Cyber Kill Chain is the “Action on Objectives”. This usually means exfiltrating data from the organisation. This may be a credit card database, it may be other key intellectual property. While Data Leakage Protection is a whole subject area in itself, my key point is that from an architecture perspective we need to think about how we prevent critical information assets from leaving an organisation. This certainly includes malicious exfiltration as part of an attack, but accidental disclosure as well.

Firstly it is essential to understand the vectors which can be used to exfiltrate information, of which there are many. For example, the Target breach used 25+ year old FTP to move a large volume on information from the organisation, apparently without detection. Is there any reason to be allowing unrestricted use of FTP through your Internet perimeter? OK so lets block it. Then its likely another exfiltration vector will be employed, possibly an SSL encrypted approach. Or, it could be uploaded via social media, or Webmail, or a Browser Based File share like Dropbox. My key point is there are a lot of these and they all need to be considered.

My recommendation is that a small set of sanctioned applications be used. Their usage should be limited to those who require them for their jobs. This is a Case-by-Case, organisation-by-organisation assessment, but it something every organisation needs to have considered and have a policy on.

Network Infrastructure - Security

At the start, I called out that network Infrastructure is a valuable target for attackers. If they have access to network infrastructure, then it is very easy to gain access to just about any resource in the network. This includes bypassing security devices such as firewalls or IPSs, or disabling network based security controls. It also becomes very easy to route traffic though alternate paths or to capture devices for interception.

All the network vendors provide comprehensive documentation on how their devices and infrastructure should be secured. These includes;

- Securing the network infrastructure itself.
- Preventing many network based attacks which can manipulate network infrastructure.

These can include attacks like ARP Spoofing, Spoofing DHCP Responses, CAM Table overflows, etc.

These need to be understood and implemented as part of network deployment.

Out of Band Management

I discussed the concept and need for Out-for-Band Management in the first White Paper. So, I'm not going to spend much time on it here, apart from to say that in this day and age, no one, should be using In-Band management. As I discussed at the start, the network infrastructure, including network security infrastructure, is a highly valuable target in itself which requires the highest degree of protection. Out-of-Band Management is neither complex nor expensive to establish and can provide a very significant level of protection from an attack trying to gain access to network infrastructure.

As such, I believe the use of OOB is an essential architectural element.

Monitoring Infrastructure

Previously I discussed the use of packet capture technologies. Today there are an increasing number of tools and Analytics solutions which will consume raw traffic streams from strategic PINs. Traditionally, traffic mirroring features such as Cisco's SPAN, or Junipers Port Mirroring provide this function. These have been a foundation technology for many years now and generally work well. However with more and more devices within a network requiring visibility of raw traffic information, the scalability of these technologies has also become an

issue in larger scale deployments. Additionally, these capabilities within virtual environments have often been very limited or lacking.

In these environments, dedicated monitoring networks have increased in prevalence. Typically a monitoring network will use a network of fibre Tap on key links in the network, to provide a large scale replication of the raw traffic stream/streams to multiple monitoring devices which need consume this traffic. The recognised industry leader in this domain is Gigamon (www.gigamon.com).

Operational Enablement

Overview

In the past, the approach taken to securing an organisation was designing a secure network which prevented intrusions or compromises, or in other words "keeping the bad guys out". This approach should absolutely remain a key goal through solid architectural principals and operational procedures.

Unfortunately, experience has shown that even with this in place, the probability of an attack successfully compromising an organisation at some point, is a high. This is one of the fundamental tenants of the Lockheed Martin Cyber Kill Chain.

We must now work on the basis that breaches are a reality. As such, a solid underlying security architecture must be augmented with;

- A capability to detect sophisticated classes of attacks.
- An ability to quickly invoke an effective incident response.
- An assumption that the use of Day-Zero malware will continue to increase and that traditional signature based systems can not be relied upon.
- Constantly gathering sufficient information to understand the damage (should it occur) and facilitate remediation.

In the past, attacks on organisations have often been considered single events. For example, a discreet event which can be blocked by an IPS. Serious attacks such as APTs are now far more complex, focused, persistent, multi-stage processes. As such, they now need to be thought of as "Campaigns" where the attackers will use multiple attack techniques, on multiple vectors, over an extended time period.

However, and to the defenders advantage, experience has shown that often the attackers will reuse tools or techniques from their tool box on multiple occasions throughout the campaign. These are known as "Indicators" and include Indicators of Attack and Indicators of Compromise. Examples of Indicators could include;

- A source IP address from which an attack is launched.
- A common theme across multiple phishing emails.
- A common piece of malware within phishing emails.
- A common group within an organisation receiving phishing emails, for example staff attending an industry conference.
- The same Day Zero, or near Day Zero malware in multiple attempts.
- Repeated usage of the same exploit or attack tool.

These Indicators need to be identified as part of security operations, seen as valuable information and used as key elements of the Detection and Incident Response processes.

Analytics

Analytics technology will be a key operational enabler going forward. Given the physical speed of networks today, coupled with instrumentation and logs from multiple sources, the ability to analyse vast amounts of data quickly to find potential security incidents, is now a fundamentally required capability. The goal of Security Analytics is to enable attack/intrusion detection as quickly as possible. This capability then enables security operations to block or stop the attack. In other words, reducing the "Dwell Time", meaning the time the attack is in progress without detection. Assuming the attack is successfully detected and stopped, Analytics can subsequently be used to provide detailed information on the attack, including its trajectory and other systems which may have been affected. Reconstruction of an attack can be invaluable for forensics and also postmortem activities.

Security Analytics goes beyond the capabilities of traditional Security Information and Event Management (SIEM) systems by incorporating a wider range of data inputs and providing far deeper analysis. It can be very challenging to detect Cyber attacks as relevant data indicating an attack is often spread across multiple devices, including servers, firewalls, IPSs, application logs and endpoints. As a result, these systems should correlate events occurring on multiple different platforms to detect suspicious patterns of activity, that would previously have gone unnoticed.

In advanced attacks, experience is showing that attackers are often very skilled at keeping their behavior patterns within the 'bell curve' of what looks normal. Identifying these level of attacks does require a new level of capability, especially in large environments.

There has been significant evidence presented from multiple sources to show that attacks can happen quickly and in a very high percentage of cases, are not detected for extended periods and often not detected by the compromised organisation themselves. Additionally, following an intrusion, organisations don't often know what data has been stolen or what other damage may have been done. Or, in many cases, even if the intrusion is still active.

Security Analytics is often coupled with threat feeds. Threat feeds provide global security intelligence, such as known malware sites, compromised devices, file hashes of known malware, and other indicators. Threat feeds are currently available from a number of vendors.

Making use of what you have

For many organisations, particularly smaller ones, budgets may simply not be available for the purchase of an Analytics platform and the associated deployment costs. In these situations, taking full advantage of 'what you already have' coupled with Open Source Tools may be the approach you need to adopt.

Throughout these White Papers I have tried to highlight many underlying areas of instrumentation. DNS and Netflow are two of the most valuable and are often completely overlooked. Open source tools like NFsen and Simple Event Correlator (SEC) can be good low cost starting points.

There is a lot which can be done with a combination of readily available Network Instrumentation and Open Source tools. Doing something, and more importantly, working it into operational processes is a far better option than doing nothing and flying blind.

Workflows

As the field of Security Analytics is relatively new, the Integration of these systems into Security Operations process and workflows is, generally speaking, in the early phases of maturity. In a later document I intend to explore this whole area further. As the focus of this Document is Security Architecture, I would like to keep the comments to ways the Analytics capability should enable a workflow or process.

I believe one of the most key considerations going forward, will be the need for a new mindset around security operations and process. We know that in the vast majority of cases, security incidents are not discovered by the affected organisation itself. For Security Operations to be effective, the thinking needs to shift to 'Active Exploration' mindset. In other words, don't wait for an alert to be raised, but using tools to actively search for anomalies, Indicators, or other interesting behavior. When something interesting is uncovered, then seeking to understand or explain.

With that said, the capabilities of the Analytics system will have a large bearing on how the workflow or process operates. In terms of capability, what we don't want is a system which just provides a long list ‘red alerts’. The goal is to have an analytics system provide correlation to a level where actionable information is the output.

Security analytics tools should offer a single point of access to event data. A consolidated view is operationally valuable as it allows Security Operations staff to zero in on the problem with the need to access multiple systems. Features such as timeline reconstruction and chronological drilldown are very useful for attack analysis and forensic investigation.

Finally I would recommend setting reasonable expectations. Going out and spending a whole lot of money on an analytics solution, thinking it will solve all your problems, will likely not achieve the overall goals. It is an investment not just in an Analytics tool set, but a companion investment is required in the integration into the operational process. This part will likely be as significant an investment as the capital cost.

Incident response - Capability to work back in time

As attack techniques continue to evolve, we must work on the assumptions that covert techniques will continue to mature and that attackers may gain a foothold without detection. Having detailed historical records of network traffic can provide a capability to work back in time and search for previous indicators of a newly discovered attack. This approach provides the capability to identify potentially compromised systems, which can then be isolated from the network and remediated.

Previously I described the use of Full-Flow Netflow as an instrumentation technique. The approach provides details of conversations occurring at specific Points-In-the-Network. The storage of this flow information, coupled with a suitable analytics system, would be one means of "working back in time" to identify previous damage from a more recently detected attack or campaign. In addition to Netflow, some Analytics systems are also utilising full packet capture information to provide richer capability, albeit at a higher cost point.

While these approaches could represent a significant investment, particularly in terms of storage. In environments where high value information assets exist, then it may well be an entirely cost justified approach.

In summary, security design needs to think along the lines that things may be missed despite considerable effort. Providing a capability to analyse and clean up past damage if detected at a latter stage, is a prudent design approach.

Wrap Up and What's Next..

I have produced two White Paper, the first explaining the fundamentals, which are still vitally important today and this second one looking at the changes.

Many aspects of IT Security today (2015) are just not working effectively. I have attempted to outline the key architectural considerations that should be incorporated going forward.

My key point from this White Paper is that many things have changed in recent years, most significantly IT Macro Trends as well as the threat landscape. These changes mean that Security and Network Security Architecture must evolve.

The threat landscape has changed a lot in the last five years and a huge amount in the last two. It will continue to change. Malware, threats and attack techniques will continue to grow in sophistication and complexity, possibly in unpredictable ways. On this basis, security design needs to think along the lines of providing capability to get ahead of the sophistication growth curve.